ProLiant Cluster HA/F500 for MA8000

Administrator Guide

Part Number 355220-006

August 2002 (Sixth Edition)

This guide provides information about the installation, configuration, and implementation of the ProLiant Cluster HA/F500 for RA8000/MA8000/ESA12000/EMA12000/EMA16000.

COMPAQ

ProLiant Cluster HA/F500 for MA8000 Administrator Guide

Contents

Chapter 2
Setting Up the ProLiant Cluster HA/F500 with Hubs

Chapter 3
Setting Up the ProLiant Cluster HA/F500 with Switches

Chapter 4
Secure Path Installation Procedures

Chapter 5
Managing the ProLiant Cluster HA/F500

Chapter 6
Troubleshooting the ProLiant Cluster HA/F500

Appendix A
Versions of Referenced Software

Glossary

Index

About This Guide

This guide provides information about the installation, configuration, and implementation of the ProLiant Cluster HA/F500 for RA8000/MA8000/ESA12000/EMA12000/EMA16000.

Intended Audience

This guide is designed for network administrators, installation technicians, systems integrators, and other technical personnel in the enterprise environment whose jobs include installing, configuring, and maintaining ProLiant clusters.

Symbols in Text

These symbols may be found in the text of this guide. They have the following meanings.

> ⚠ **WARNING: Text set off in this manner indicates that failure to follow directions in the warning could result in bodily harm or loss of life.**

> △ **CAUTION:** Text set off in this manner indicates that failure to follow directions could result in damage to equipment or loss of information.

IMPORTANT: Text set off in this manner presents clarifying information or specific instructions.

NOTE: Text set off in this manner presents commentary, sidelights, or interesting points of information.

Text Conventions

This document uses the following conventions:

- *Italic type* is used for complete titles of published guides or variables. Variables include information that varies in system output, in command lines, and in command parameters in text.

- **Bold type** is used for emphasis, for onscreen interface components (window titles, menu names and selections, button and icon names, and so on), and for keyboard keys.

- `Monospace typeface` is used for command lines, code examples, screen displays, error messages, and user input.

- Sans serif typeface is used for uniform resource locators (URLs).

Getting Help

If you have a problem and have exhausted the information in this guide, you can get further information and other help in the following locations.

Compaq Technical Support

In North America, call the Compaq Technical Support Phone Center at 1-800-652-6672. This service is available 24 hours a day, 7 days a week. For continuous quality improvement, calls may be recorded or monitored. Outside North America, call the nearest Compaq Technical Support Phone Center. For telephone numbers of worldwide Technical Support Centers, go to www.compaq.com.

Be sure to have the following information available before you call Compaq:

- Technical support registration number (if applicable)

- Product serial number

- Product model name and number

- Applicable error messages

- Add-on boards or hardware

- Third-party hardware or software

- Operating system type and revision level

Compaq Website

For information on this product, as well as the latest drivers and flash ROM images, go to www.compaq.com.

Compaq Authorized Reseller

For the name of your nearest Compaq authorized reseller:

- In the United States, call 1-800-345-1518.

- In Canada, call 1-800-263-5868.

- Elsewhere, see the Compaq website for locations and telephone numbers.

Reader's Comments

Compaq welcomes your comments on this guide. Please send your comments and suggestions by email to ServerDocumentation@compaq.com.

1

Overview of the ProLiant Cluster HA/F500

The *ProLiant*™ Cluster HA/F500 is a two-, three-, or four-node cluster solution composed of Compaq ProLiant servers and *StorageWorks*™ storage components by Compaq; it executes on a Microsoft Windows NT Server 4.0, Enterprise Edition platform with Microsoft Cluster Server (MSCS), Microsoft Windows 2000 Advanced Server with Microsoft Cluster Service, or Microsoft Windows 2000 Datacenter Server with Microsoft Cluster Service. The ProLiant Cluster HA/F500 provides the highest levels of availability and can be configured to provide an environment in which there are no single points of failure (NSPOF), creating a cluster environment ideal for business-critical applications and data access requirements.

The "ProLiant Cluster HA/F500 Configurations" section of this chapter describes supported configurations for the HA/F500. These configurations include:

- Basic HA/F500 configurations with hubs

- Enhanced HA/F500 configuration with hubs

- Multiple cluster configurations with hubs

- Basic HA/F500 configurations with switches

- Enhanced HA/F500 configuration with switches

- Microsoft Datacenter HA/F500 configuration with switches

- Multiple cluster configurations with switches

The ProLiant Cluster HA/F500 is composed of a number of different hardware and software products. The "Components" section of this chapter discusses how each of these products plays a role in bringing a complete clustering solution to your computing environment.

- Compaq ProLiant servers

- Compaq Fibre Channel storage hubs or Compaq Fibre Channel storage switches

- Host bus adapters

- Storage controllers

- Storage subsystem

- Cluster interconnect

- Microsoft Windows NT Server 4.0, Enterprise Edition:
  — Microsoft Cluster Server (MSCS)
  — Microsoft Cluster Administrator

- Microsoft Windows 2000 Advanced Server:
  — Microsoft Cluster Service
  — Microsoft Cluster Administrator

- Microsoft Windows 2000 Datacenter Server:
  — Microsoft Cluster Service
  — Microsoft Cluster Administrator

- Compaq Software:
  — Compaq *SmartStart*™ CD
  — Compaq StorageWorks Command Console
  — *SANworks*™ Secure Path by Compaq for Windows NT and Windows 2000
  — *Compaq Insight Manager*™ 7

ProLiant Cluster HA/F500 Configurations

The ProLiant Cluster HA/F500 configurations can involve either hubs or switches. The following text is divided into two sections: a discussion of configurations using hubs, followed by a discussion of configurations using switches.

The ProLiant Cluster HA/F500 configurations support both Fibre Channel hubs and Fibre Channel switches with the Array Controller Software (ACS) version 8.6. ACS 8.7 will only support Fibre Channel switches.

The ProLiant Cluster HA/F500 supports Disaster Tolerant (DT) configurations. Refer to the Compaq StorageWorks website for DT ordering and configuration information at www.compaq.com/highavailability.

ProLiant Cluster HA/F500 Configurations with Hubs

The ProLiant Cluster HA/F500 can be set up in a number of different configurations involving servers, hubs, and storage subsystems connected through a Fibre Channel Arbitrated Loop (FC-AL).

- The basic configuration is a cluster with a single host bus adapter in each server, one hub, and one or two storage controllers. The enhanced configuration is a cluster with two host bus adapters in each server, two hubs, and two storage controllers, giving it increased availability over the basic configurations.

- Additionally, two clusters can be configured to use the same storage subsystems.

In any cluster, a maximum of two storage controller pairs can be connected to a single loop. This limitation dictates how many storage subsystems can be used in the cluster (a maximum of two storage subsystems).

The servers in a ProLiant Cluster HA/F500 are also connected to a local area network (LAN), usually referred to as the client LAN, through which users access the cluster. There is also an interconnect connecting the two servers, frequently referred to as the cluster interconnect. The cluster interconnect allows the two servers to communicate without the delay potentials inherent on the client LAN. Because the client LAN and the cluster interconnect have similar configurations, they are not mentioned any further when discussing the different configurations in the following sections.

Basic Configurations with Hubs

The basic configurations with hubs are shown in Figure 1-1 and Figure 1-2. Each of the configurations contains a single host bus adapter in each server and a single hub.

Single Storage Controller

Figure 1-1 shows a configuration with a single connection between the hub and the storage controller. Failures in the server, server software, host bus adapter, or cable from the host bus adapter to the hub will cause Microsoft cluster software failover to the second server, and the cluster will continue to service data requests from the second server.

NOTE: The dashed lines indicate maximum configuration.
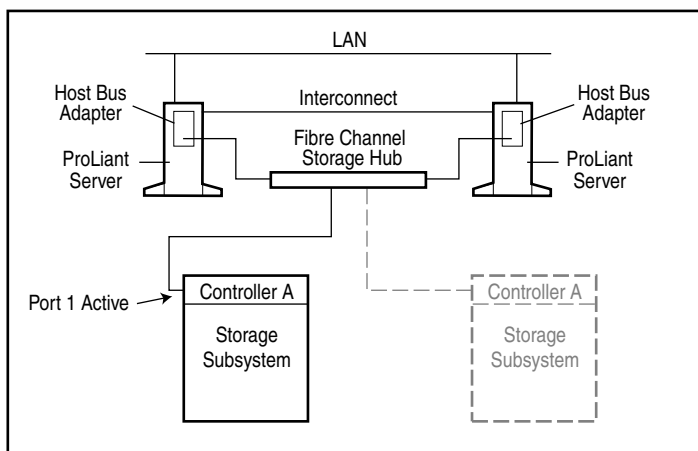
Figure 1-1: Basic configuration with a hub (single storage controller)

Dual Storage Controllers

Figure 1-2 adds a second storage controller and a second cable between the hub and the additional storage controller. This configuration results in a higher level of availability because of the additional redundant components.

NOTE: The dashed lines indicate maximum configuration.
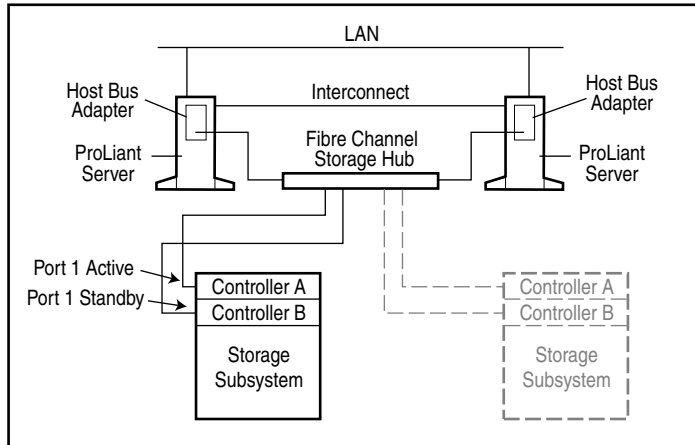
Figure 1-2: Basic configuration with a hub (dual storage controllers)

Enhanced Configuration with Hubs

Figure 1-3 shows a configuration without any single points of failure. It improves on the basic configuration shown (see Figure 1-2) by adding a second host bus adapter in each server and a second hub. The combination of second adapter, hub, and controller form a second independent path to the storage subsystem.

To allow dual paths to the storage, the Secure Path software must be installed on both servers. With Secure Path, data can flow simultaneously over both host bus adapters to the storage subsystem, and static load balancing can be performed over the two paths to help maximize performance.

NOTE: The dashed lines indicate maximum configuration.
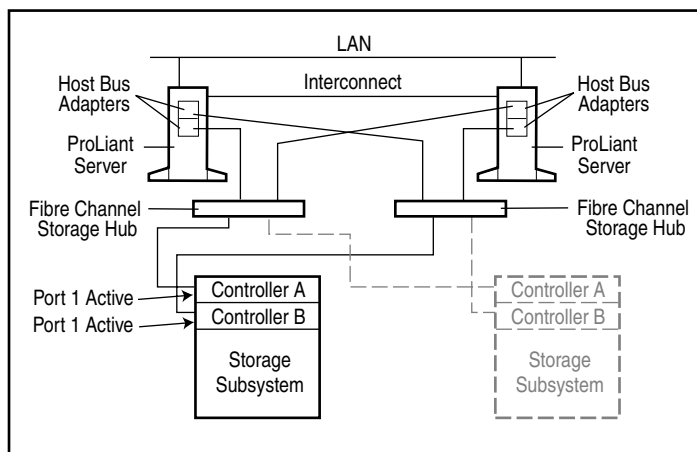
Figure 1-3: Enhanced configuration with hubs

A component failure in this cluster will result in a failover to a second component, and the end user can continue using the cluster. Some typical failures and responses in the enhanced configuration include:

- A server failure will cause Microsoft cluster software to fail over to the second server.

- A host bus adapter failure will cause subsequent data requests intended for the failed adapter to be routed over the remaining good adapter.

- A hub or cable failure will be detected as a host bus adapter failure, and a failover to the second adapter, which is using the remaining good hub and good cables, will occur.

- A controller failure will cause the second controller to take over for the failed controller. Secure Path will then route the data requests to the second controller.

In all of the previous cases, interruptions to the end user are minimal and in some cases may not even be noticeable.

Multiple Cluster Configurations with Hubs

Any two of the same basic or enhanced clusters with hubs can be combined into a single configuration with both clusters accessing the same storage subsystems. A sample configuration is shown in Figure 1-4. Note in particular that the two clusters cannot share the same hub and that one cluster is always attached to port 1 on the controllers and the other cluster is always attached to port 2 on the controllers.
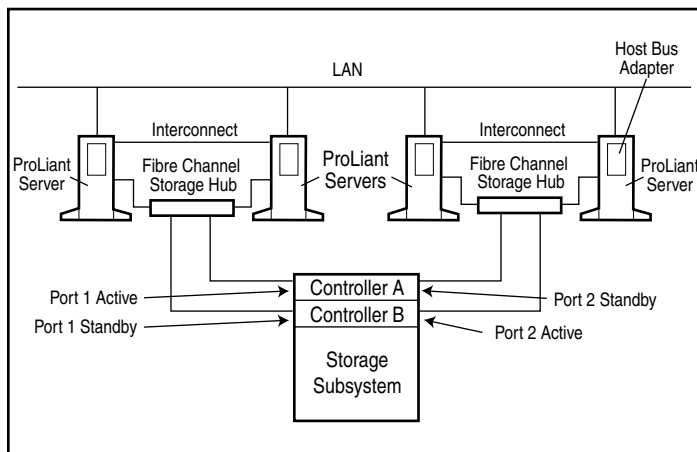
Figure 1-4: Multiple cluster configuration with hubs

ProLiant Cluster HA/F500 Configurations with Switches

The ProLiant Cluster HA/F500 can be set up in a number of different configurations involving servers, switches, and storage subsystems connected through a Fibre Channel switch fabric (FC-SW).

- The basic configuration is a cluster with a single host bus adapter in each server, one switch, and one or two storage controllers.

- The enhanced configuration is a cluster with two host bus adapters in each server, two switches, and two storage controllers, giving it increased availability over the basic configurations.

- Additionally, two to four clusters can be configured to use the same storage subsystems.

The basic and enhanced configurations with switches are discussed in the following sections. In an HA/F500 configuration, a maximum of eight storage controller pairs can be connected to a single switch. This limitation dictates how many storage subsystems can be used in the cluster (a maximum of eight storage subsystems or some combination of each type of storage unit). There can be a maximum of 672 drives in a configuration. A single cluster in an HA/F500 configuration can interact with a maximum of eight pairs of controllers.

The servers in a ProLiant Cluster HA/F500 are also connected to a LAN, usually referred to as the client LAN, through which users access the cluster. There is also an interconnect connecting the two servers, frequently referred to as the cluster interconnect. The interconnect is used to allow the two servers to communicate without the delay potentials inherent on the client LAN. Because the client LAN and the cluster interconnect are the same in all the configurations, they are not mentioned any further when discussing the different configurations.

Basic Configurations with Switches

The basic configurations with switches are shown in Figure 1-5 and Figure 1-6. Each of the configurations contains a single host bus adapter and a single switch.

Single Storage Controller

Figure 1-5 shows a configuration with a single connection between the switch and the storage controller. Failures in the server, server software, host bus adapter, or cable from the host bus adapter to the switch will cause Microsoft cluster software failover to the second server, and the cluster will continue to service data requests from the second server.

NOTE: The dashed lines indicate maximum configuration.
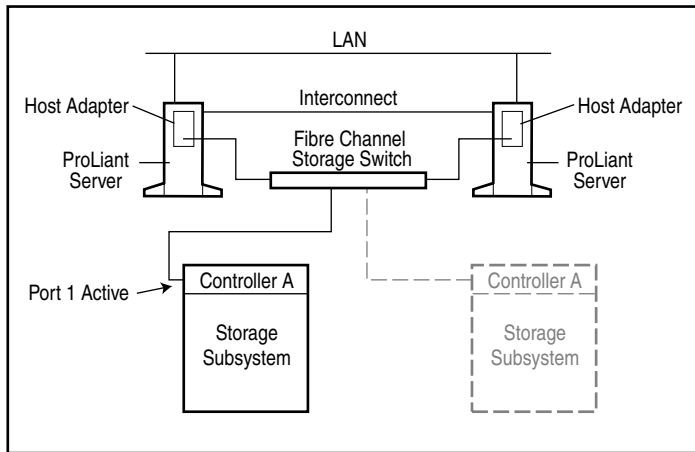
Figure 1-5: Basic configuration with a switch (single storage controller)

Dual Storage Controllers

Figure 1-6 adds a second storage controller and a second cable between the switch and the additional storage controller. This results in a higher level of availability because of the additional redundant components.

NOTE: The dashed lines indicate maximum configuration.
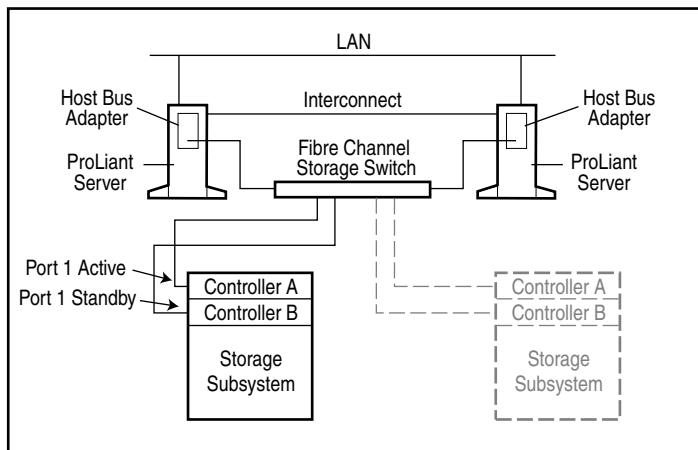
Figure 1-6: Basic configuration with a switch (dual storage controllers)

Enhanced Configuration with Switches

Figure 1-7 is a configuration without any single points of failure. It improves on the basic configuration shown in Figure 1-6 by adding a second host bus adapter to each server and a second switch. The combination of second adapter, switch, and controller form a second independent path to the storage subsystem.

To allow dual paths to the storage, the Secure Path software must be installed on all servers. With Secure Path, data can flow simultaneously over both host bus adapters to the storage subsystem, and you can perform static load balancing over the two paths to help maximize performance.

NOTE: The dashed lines indicate maximum configuration.
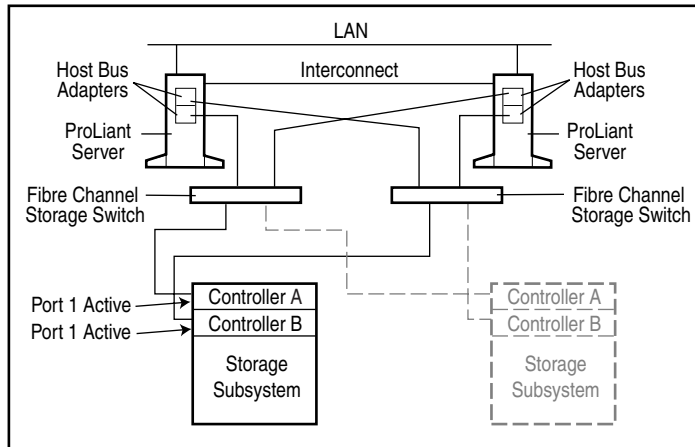
Figure 1-7: Enhanced configuration with switches

A component failure in this cluster will result in a failover to a second component, and the end user can continue using the cluster. Some typical failures and responses in the enhanced configuration include:

- A server failure will cause Microsoft cluster software to fail over to the other node.

- A host bus adapter failure will cause subsequent data requests intended for the failed adapter to be routed over the remaining good adapter.

- A switch or cable failure will be detected as a host bus adapter failure, and a failover to the second adapter, which is using the remaining good switch and good cables, will occur.

- A controller failure will cause the second controller to take over for the failed controller. Secure Path will then route the data requests to the second controller.

In all of the previous cases, interruptions to the end user are minimal and in some cases may not even be noticeable.

Microsoft Datacenter HA/F500 Configuration with Switches

Microsoft Datacenter is an enhanced configuration with switches. It improves on the enhanced configuration shown in Figure 1-7 by adding one or two more nodes. Figure 1-8 shows a four-node Microsoft Datacenter configuration.
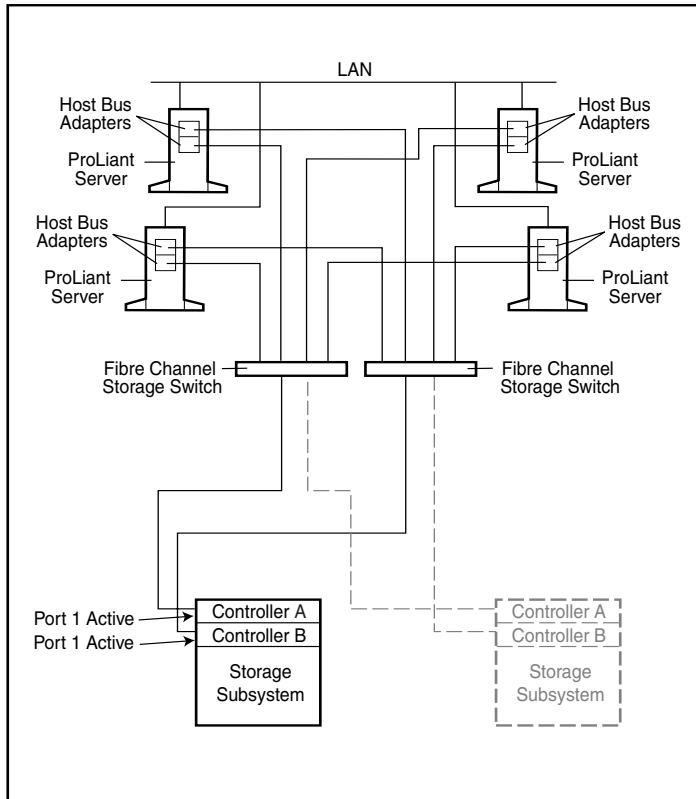
Figure 1-8: Microsoft Datacenter configuration with switches

Multiple Cluster Configurations with Switches

Up to eight clusters can be combined into a single HA/F500 configuration with the clusters accessing the same group of storage subsystems. A sample configuration with four clusters is shown in Figure 1-9. The figure shows four storage subsystems.
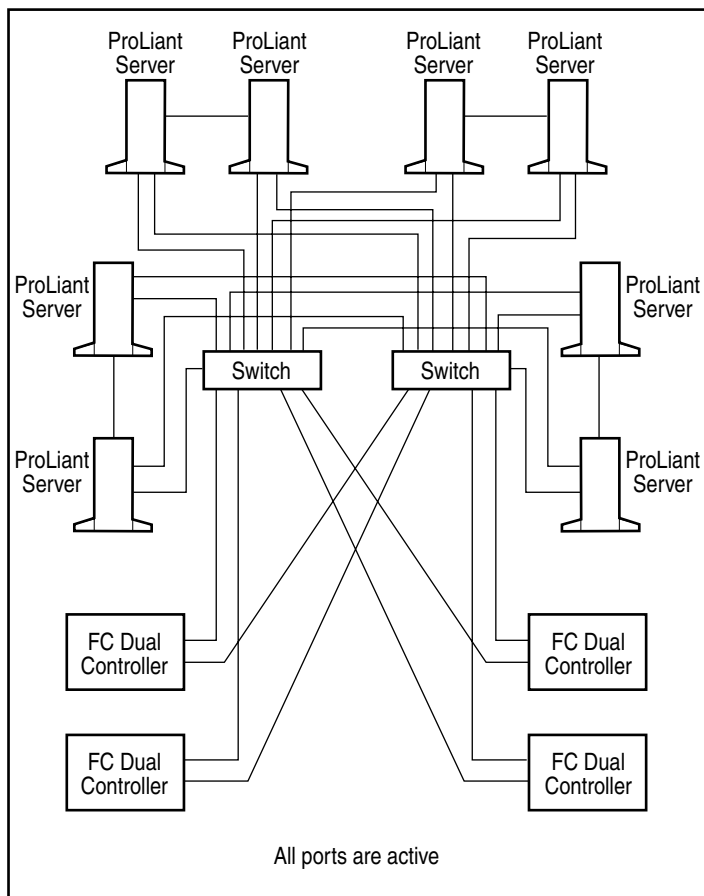
Figure 1-9: Multiple cluster configuration with switches

Hub and Switch Maximums

ProLiant Cluster HA/F500 configurations can involve either hubs or switches. Table 1-1 shows the numerical maximums for hub and switch configurations. Hub cascading is not supported. For cascading switches, refer to the StorageWorks documentation.

Table 1-1: Hub and Switch Maximums

Configurations	Hub **Single Cluster Configuration**	Hub **Multiple Cluster Configuration**	Switch **Single Cluster Configuration**	Switch **Multiple Cluster Configuration**
Storage controller pairs per host bus adapter	2	2	8	8
Hubs or switches per storage controller pair	2	2	2	2
Clusters (two servers each) per storage controller pair	1	2	1	8
Host bus adapters per host per storage controller pair	2	2	2	2
Storage controller pairs	2	2	8	8
Storage Subsystem Drives	168	168	672	672

Components

The following sections provide information on the components supported in a ProLiant Cluster HA/F500 for MA8000.

Compaq ProLiant Servers

The HA/F500 supports high-end and high-density Compaq ProLiant servers. Compaq high-end servers incorporate many high-availability and manageability features, such as an offline backup processor, PCI Hot Plug, redundant hot-plug fans, redundant processor power modules, redundant network interface card (NIC) support, hot-plug dual-port 10/100 TX NIC, and redundant hot-plug power supplies. To obtain an up-to-date list of supported combinations, go to www.compaq.com/highavailability.

Fibre Channel Storage Hubs

The ProLiant Cluster HA/F500 can use an FC-AL to communicate between the ProLiant servers and the storage subsystems. The central component of an FC-AL is the hub. The ProLiant Cluster HA/F500 allows use of either the Fibre Channel Storage Hub 7 (7 ports) or the Fibre Channel Storage Hub 12 (12 ports). If a ProLiant Cluster HA/F500 configuration requires the use of more than one hub, any combination of 7- and 12-port hubs can be used. If a configuration involves more than one cluster, each cluster must have dedicated hubs.

Compaq Fibre Channel Storage Switches

The ProLiant Cluster HA/F500 supports the use of the following Compaq Fibre Channel storage switches:

- Compaq StorageWorks Fibre Channel SAN Switch 8 (8 ports)

- Compaq StorageWorks Fibre Channel SAN Switch 16 (16 ports)

- Compaq StorageWorks Fibre Channel SAN Switch 8-EL (8 ports)

- Compaq StorageWorks Fibre Channel SAN Switch 16-EL (16 ports)

- Compaq StorageWorks Fibre Channel SAN Switch (2 Gb) 16 ports

- Compaq StorageWorks Fibre Channel SAN Switch Integrated 32

- Compaq StorageWorks Fibre Channel SAN Switch Integrated 64

- Compaq StorageWorks SAN Director 64 (1 Gb) Switch

If a ProLiant Cluster HA/F500 configuration requires the use of more than one switch, any combination of 8- and 16-port switches can be used. If a configuration involves more than one cluster, the different clusters may share the same fabric.

Host Bus Adapters

The host bus adapter is a PCI card that connects one ProLiant server to:

- The FC-AL to which the hub, storage subsystem, and second server are also connected

   or

- A Compaq Fibre Channel Switch, to which the storage subsystem and second server are also connected

Depending upon the availability requirements of the cluster, each ProLiant server can have one or two host bus adapters. The basic configurations (described previously) require one host bus adapter in each server; the enhanced configuration requires two host bus adapters and Secure Path software in each server.

Storage Controllers

The storage controller (see Figure 1-10) is a dual-port Fibre Channel RAID controller installed in the storage subsystem. It manages drives and connects the storage subsystem to hubs or switches.

The storage controllers include the following features:

- PCMCIA slot (1)

- Port 2 (2)

- Port 1 (3)

- Maintenance port (4)

- LEDs (5)
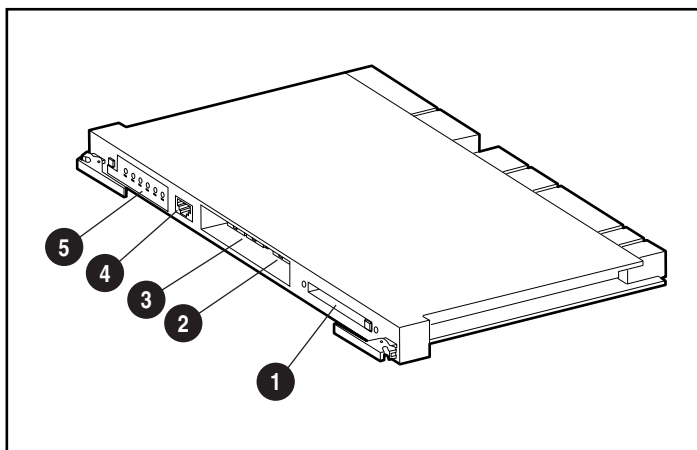
Figure 1-10: Storage controller

The storage controllers can be configured as follows:

- For the basic configuration using one controller, the single storage controller is configured as active. The storage controller is connected through a fiber optic cable to the hub or switch.

- For the basic configuration using two storage controllers, the controllers are configured in transparent mode. In transparent mode, one controller is configured as active, and the second controller is configured as standby. All data requests pass through the active controller. If the active controller fails, the standby controller becomes active and takes over the role of the failed controller. Both storage controllers are connected through fiber optic cables to different ports on the same hub or switch.

- For the enhanced configuration, both storage controllers are configured as active, but access to a logical unit (LUN) is restricted to either port 1 or port 2 to ensure data integrity. In this configuration, both storage controllers are receiving and transmitting data, increasing the data throughput for the cluster over the basic configuration. If the first controller fails, the other storage controller takes over the role of the failed controller, in addition to its own, and continues operation. Port *N* (where *N*=1 or 2) on one storage controller is connected through a fiber cable to one port on a hub or switch; port *N* on the other storage controller is connected to the second hub or switch. Connecting the storage controllers in this manner provides two independent paths to the storage subsystem.

Storage Subsystem

The storage subsystems for the ProLiant Cluster HA/F500 are the Compaq StorageWorks RAID Array 8000 (RA8000), the Compaq StorageWorks Enterprise Storage Array 12000 (ESA12000), the Compaq StorageWorks Modular Array 8000 (MA8000), the Compaq StorageWorks Enterprise Modular Array 12000 (EMA12000), and the Compaq StorageWorks Enterprise Modular Array 16000 (EMA16000).

The RA8000/MA8000 and ESA12000/EMA12000/EMA16000 are high-capacity storage subsystems that support both single and dual RAID controllers. The RA8000/MA8000 and ESA12000/EMA12000/EMA16000 support up to 72 disks per controller or controller pair.

High-availability components of the storage subsystems include redundant power supplies, fans, controllers, cache battery backup, hot-global spare drives, and multilevel RAID architecture.

For more information about the storage subsystem components, refer to:

- Documentation provided with the storage subsystem

- The StorageWorks website at www.compaq.com/storage

Cluster Interconnect

The cluster interconnect is a data path over which nodes of a cluster communicate. This type of communication is termed *intracluster communication.* At a minimum, the interconnect consists of two network adapters (one in each server) and a cable connecting the adapters.

The cluster nodes use the interconnect data path to:

- Communicate individual resource and overall cluster status

- Send and receive heartbeat signals

- Update modified registry information

IMPORTANT: TCP/IP must be used as the cluster communication protocol. When configuring the interconnects, be sure to enable TCP/IP.

Redundant Interconnects

To reduce potential disruptions of intracluster communication, connect a redundant path over which communication can continue if the primary path is disrupted.

Compaq recommends that you configure the client LAN as a backup path for intracluster communication.

Compaq offers a feature that configures two Compaq Ethernet adapters (or two ports on a single adapter) so that one is a hot-backup for the other. This feature, called NIC Teaming, is available on all Compaq 10/100 Fast Ethernet products. The Compaq NIC Teaming software is available on the Compaq Server Support (SSD) for Microsoft Windows NT 4.0, or on the Compaq Support Paq for Microsoft Windows 2000. This software can be downloaded from the Compaq website and is included with each server as part of the SmartStart utilities.

NIC Teaming should only be used with NICs that connect to the client LAN. This feature should not be used with NICs used for the dedicated intracluster communication link. For detailed information about interconnect redundancy, refer to the Compaq white paper, *Increasing Availability of Cluster Communications in a Windows NT Cluster* which is available at www.compaq.com/highavailability.

Interconnect Adapters

Ethernet adapters and switches are supported as interconnects in ProLiant clusters. Either 10-Mb/sec or 100-Mb/sec Ethernet may be used.

NOTE: For a list of supported interconnect adapters, see the Microsoft Windows NT, Windows 2000 Advanced Server, Windows 2000 Datacenter Server, and Microsoft cluster software compatibility list available from the Microsoft website at www.microsoft.com. Be sure that the adapter you choose is on the list.

The Ethernet adapters can be physically connected with an Ethernet crossover cable (direct connection) or through an Ethernet hub.

NOTE: An Ethernet crossover cable is provided in the HA/F500 Basic and Enhanced Cluster Kits. For the Microsoft Datacenter configuration using more than two nodes, an Ethernet hub is required.

Ethernet Direct Connection

A direct Ethernet connection uses only three components:

- Two interconnect adapters

- One Ethernet crossover cable

This solution is inexpensive, easy to install, and easy to troubleshoot. It also minimizes the potential single points of failure.

Connecting interconnect adapters directly to each other requires a special cable. If Ethernet is used, an Ethernet crossover cable (included in the ProLiant Cluster HA/F500 Cluster Kit) must be used.

If the Ethernet Direct Connection cable supplied with your kit is used and Windows 2000 is installed, the interconnect network may not show up during the cluster installation. This is because the connection is only shown if it is currently active at the time of installation. If the other cluster nodes are powered off at the time you install MSCS, the connection is considered by Windows 2000 to be inactive. In this case, define the existing public network connection as "all communications" during the installation. After MSCS is configured on all nodes, the interconnect network should automatically appear in the networks group in Cluster Administrator.

Use the following steps to configure the networks for MSCS use after installing Windows 2000:

1. Right-click the cluster name in Cluster Administrator.

2. Select **Properties.**

3. Choose the **Network Priority** tab from the dialog box.

4. Configure the network roles as necessary.

Microsoft Datacenter Ethernet Connection

The Microsoft Datacenter configuration uses Ethernet hubs and cables for the dedicated interconnect. A Microsoft Datacenter Ethernet connection uses the following components:

- Four interconnect adapters

- Four Ethernet cables

- One Ethernet hub

Local Area Network

Every client/server application requires a LAN, over which client machines and servers communicate. The components of the LAN on a cluster are no different than with a stand-alone server configuration.

Since clients desiring the full advantage of the cluster will now connect to the cluster rather than to a specific server, configuring client connections will differ from those for a stand-alone server. Clients will connect to virtual servers that are cluster groups containing their own IP addresses.

2

Setting Up the ProLiant Cluster HA/F500 with Hubs

Preinstallation Instructions

Before setting up the ProLiant Cluster HA/F500 with hubs, verify that the hardware and software kits are appropriate for this installation. The following section gives more details about the different installation configurations.

IMPORTANT: ACS 8.6 supports the use of hubs. ACS 8.7 does not support the use of hubs.

Checking the Hardware and Software

Use the following tables to verify that the appropriate hardware and software installation kits were received. In the following tables, the numbers in parentheses refer to configurations with two storage subsystems.

Table 2-1: Basic Configuration with Hubs (Single Storage Controller)

Quantity	Description
2	ProLiant servers
2	Host bus adapters
1 (2)	Storage subsystem with one storage controller per subsystem

continued

Table 2-1: Basic Configuration with Hubs (Single Storage Controller)
continued

Quantity	Description
1	Storage controller platform kit
1 (2)	ACS Controller software
1	7- or 12-port Fibre Channel hub
1	ProLiant Cluster HA/F500 Basic Kit with Ethernet crossover cable for the server interconnect
1	Microsoft operating system, which includes Microsoft cluster software
As required	Optical Fibre Channel interconnect hardware

Table 2-2: Basic Configuration with Hubs (Dual Storage Controller)

Quantity	Description
2	ProLiant servers
2	Host bus adapters
1 (2)	Storage subsystem with two storage controllers per subsystem
1	Storage controller platform kit
2 (4)	ACS Controller software
1	7- or 12-port Fibre Channel hub
1	ProLiant Cluster HA/F500 Basic Kit with Ethernet crossover cable for the server interconnect
1	Microsoft operating system, which includes Microsoft cluster software
As required	Optical Fibre Channel interconnect hardware

Table 2-3: Enhanced Configuration with Hubs

Quantity	Description
2	ProLiant servers
4	Host bus adapters
1 (2)	Storage subsytem with two storage controllers per subsystem
1	Storage controller platform kit
2 (4)	ACS Controller software
2	7- or 12-port Fibre Channel hubs
1	ProLiant Cluster HA/F500 Enhanced Kit containing an Ethernet crossover cable for the server interconnect and Secure Path software
1	Microsoft operating system, which includes Microsoft cluster software
As required	Optical Fibre Channel interconnect hardware

If any required component is missing, contact your local sales representative or call 1-800-652-6672.

Installing the Hardware

The following sections provide information on installing the hardware components used in a ProLiant cluster.

Setting Up the Servers

Follow the installation instructions in the ProLiant server documentation to set up each server as a stand-alone server. Then use the instructions in the following sections to configure the cluster.

Setting Up the Storage Subsystem

Refer to the documentation that was shipped with the storage subsystem for detailed installation instructions.

Installing the Disk Drives, Storage Controller, and Program Card

The storage subsystem uses hard drives, storage controllers, and PCMCIA program cards that must be installed (see Figure 2-1).

IMPORTANT: If you are setting up a basic configuration, verify that only one controller is installed. If more than one controller is present, remove the additional controllers to facilitate installation of the HA/F500.

To install a disk drive (1):

1. Insert a disk drive into the shelf guide slots.

2. Slide the disk drive into the shelf and use the attached lever to push it into place.

IMPORTANT: For optimum SCSI bus distribution, install the disk drives along different SCSI buses.

To install the storage controller (2):

1. Insert the storage controller into the top shelf guide slots.

2. Slide the storage controller into the shelf until the mounting tabs snap into place.

To install the PCMCIA program card (3) in the storage controller:

1. Remove the program card cover (4) from the controller PCMCIA slot.

2. Insert the PCMCIA program card (3) into the top controller slot.

3. Replace the program card cover (4) over the controller slot.
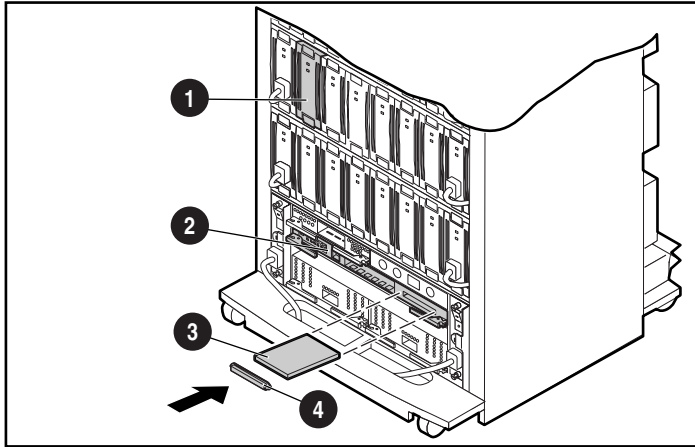
Figure 2-1: Installing the disk drives, storage controller, and PCMCIA cards (RA8000 shown)

Installing the Host Bus Adapter

Follow the installation instructions in the server documentation to install the host bus adapter in the servers. Install one adapter in each server.

For the basic configurations (single host bus adapter, single or dual storage controllers), only one adapter is installed per server. Chapter 4 discusses the enhanced configuration (dual host bus adapters, dual storage controllers) in which a second adapter is installed.

Connect the adapter from each server to the Fibre Channel hub. Connect port 1 (the left port) of the storage controller to the Fibre Channel hub.

Designating the Server as a Maintenance Terminal

A server must be connected to the storage controller to provide a maintenance terminal.

NOTE: Only one server should be designated as the maintenance server. It is recommended that a separate stand-alone server that is not part of the cluster be designated as the maintenance server.

To connect the server to the storage controller (see Figure 2-2):

1. Connect the RJ-12 connector on the communications cable to the maintenance port (1) on the storage controller.

2. Connect the 9-pin serial connector on the communications cable to either the COM1 or COM2 port (2) on the server.

NOTE: Be careful to note which serial port is used. This information will be needed when setting up the communications program and configuring the controller.
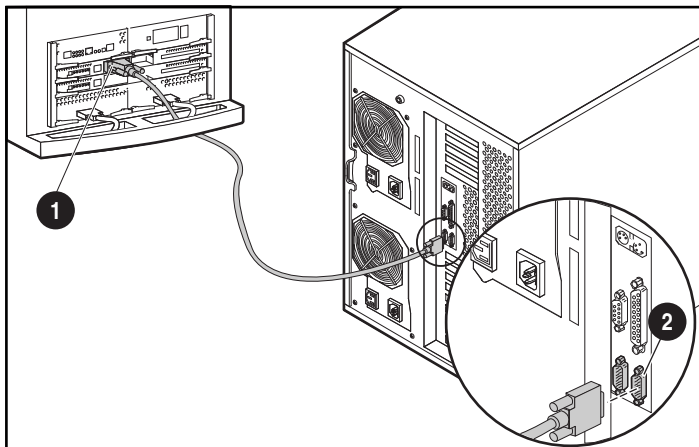
Figure 2-2: Designating the maintenance terminal

Turning On the Subsystem Power

To power on the subsystem (see Figure 2-3):

1. Connect the storage subsystem cabinet to an AC power outlet (1).

2. Turn the storage subsystem power to the on position (2). (This refers to RA8000/ESA12000 storage subsystems.)

3. Wait until the storage subsystem is completely booted and ready to operate.

4. Turn on both servers.

The storage subsystem is ready to operate when the Reset LED on the storage controller operates at a rate of one blink per second.
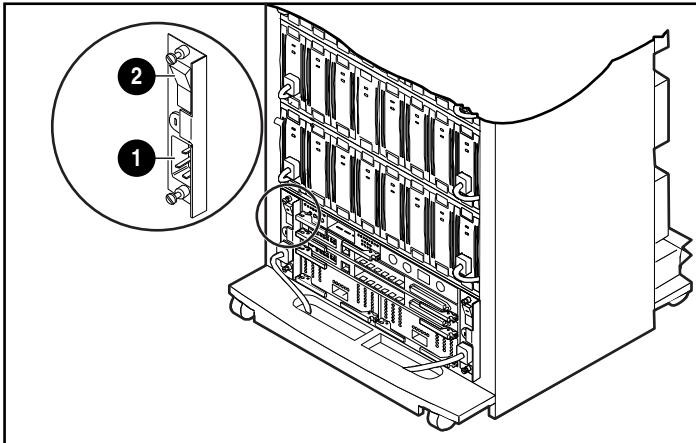
Figure 2-3: Turning on the subsystem power (RA8000/ESA12000)

Installing the StorageWorks Command Console Client

Install the StorageWorks Command Console (SWCC) Client software on the host designated as the maintenance terminal.

To install the SWCC Client:

1. Insert the storage subsystem CD-ROM from the HSG80 NT Intel platform kit in the CD-ROM drive.

2. Select **Solution Software for ACS.**

3. Click **Skip Driver Upgrade.**

4. Select **StorageWorks Command Console (SWCC).**

5. Select **CLI**, **Command Window, HSG80 or Newer.**

6. Follow the instructions on the screen.

Configuring the Storage Subsystem for Fibre Channel

The storage subsystem is preconfigured with specific parameters. These parameters may need to be changed depending on which cluster configuration is chosen. This section reviews the required steps to make these changes.

Setting Controller to No Failover

To set the storage controller to no failover:

1. Click **Start** on the taskbar.

2. Select **Programs, Command Console**, **CLI Window** to display the **Connection Selection** dialog box.

3. Select **Serial** and click **OK** to display the **Connect Serial** dialog box.

4. Select the COM port (the port to which the storage controller is connected).

5. Select a baud rate of **9600.**

6. Click **Connect** to display the CLI window.

> **NOTE:** You may need to press **Enter** to initialize communications before entering CLI commands.

7. Type the following command and press **Enter:**

   ```
   show this_controller
   ```

8. Verify that the "Not Configured for Dual-redundancy" text is displayed in the controller properties.

9. If not, type the following command and press **Enter:**

   ```
   set nofailover
   ```

> **NOTE:** If you are unable to connect with the storage controller, reset the controller by pressing the flashing Reset LED button on the controller, and retry the connection after two minutes.

Disabling the Command Console LUN

> **NOTE:** The Command Console LUN is disabled by default for SCSI mode SCSI-2. For SCSI mode SCSI-3, the Command Console LUN cannot be disabled.

To disable the Command Console LUN:

1. Click **Start** on the Windows taskbar.

2. Select **Programs, Command Console, CLI Window** to display the **Connection Selection** dialog box.

3. Select **Serial** and click **OK** to display the **Connect Serial** dialog box.

4. Select the COM port (the port to which the storage controller is connected).

5. Select a baud rate of **9600.**

6. Click **Connect** to display the CLI window.

7. Type the following command and press **Enter:**

   ```
   set this_controller nocommand_console_lun
   ```

8. After the prompt returns, type the following and press **Enter:**

   ```
   show this_controller
   ```

9. Verify that the "Command Console LUN disabled" text appears in the controller properties. If so, then continue with the next section.

Setting Port Topologies and the Arbitrated Loop Physical Address

Depending on the configuration, you may need to change the port topology and the Arbitrated Loop Physical Address (ALPA) settings. To set port topology and ALPA settings:

NOTE: Refer to the HSG80 documentation for detailed instructions on changing the port topology and ALPA settings.

1. From the CLI window, type the following and press **Enter:**

   ```
   set this_controller port_1_al_pa=71
   ```

2. Type the following and press **Enter:**

   ```
   set this_controller port_1_topology=loop_hard
   ```

3. Type the following and press **Enter:**

   ```
   set this_controller port_2_topology=offline
   ```

4. Restart the controller by typing the following and pressing **Enter:**

   ```
   restart this_controller
   ```

5. Wait two minutes for the controller to restart.

6. Type the following and press **Enter:**

   ```
   show this_controller
   ```

7. Verify the following:

   — port_1_topology = loop hard (Loop Up)

   — port_1_al_pa = 71 (71 Negotiated)

   — port_2_topology = OFFLINE (Offline)

Discovering Devices

To enable the controller to discover the disk drives:

1. From the CLI window, type the following and press **Enter:**

   run config

2. Wait a few minutes for the controller to scan and discover the disk drives.

3. Type the following and press **Enter:**

   show devices

4. Verify that the devices discovered are consistent with the disk drives installed.

   The disk drives will be named in the DISK##### format. Refer to the storage subsystem documentation for the naming conventions for storage subsystem devices.

5. Close the CLI window.

Installing the Adapter Driver

To install the adapter driver:

IMPORTANT: If the host bus adapter drivers are being installed on a new system with a Windows NT operating system, you need to make sure an older version of the driver is installed on the system first to prevent an error with the upgrade utility. An older version of the host bus adapter driver can be downloaded from www.compaq.com/products/storageworks/adapters.html.
A new system with a Windows 2000 operating system does not require a previous driver version.

1. Insert the storage subsystem CD-ROM from the HSG80 NT Intel platform kit in the CD-ROM drive.

2. Select **Solution Software for ACS.**

3. Select **KGPSA Driver Upgrade.**

4. Follow the instructions on the screen.

5. Repeat steps 1 through 4 on the other servers.

6. Verify the HBA driver version installed on the server.

7. Go to www.compaq.com/products/storageworks/adapters.html and check for the latest Fibre Channel host bus adapter driver version supported.

   IMPORTANT: Skip the following step if you have the latest Fibre Channel HBA driver installed on the server.

8. Download the latest Fibre Channel HBA driver utility and install the utility on each server, one server at a time.

Running the FC-AL Setup Utility

IMPORTANT: Run this utility before any NT File System (NTFS) volumes are created on the storage subsystem. If NTFS volumes exist, running FC-AL Setup and rebooting may reassign drive letters. Refer to the storage subsystem documentation for details on drive letter reassignment.

The FC-AL Setup Utility installs the version of the HSZDISK driver required for Windows NT Clusters and sets the ALPA and registry parameters.

To run the FC-AL Setup Utility:

1. Insert the storage subsystem CD-ROM from the HSG80 NT Intel platform kit in the CD-ROM drive.

2. Select **Solution Software for ACS.**

3. Click **Skip Driver Upgrade.**

4. Select **Fibre Channel Software Setup.**

5. Select one of the Fibre Channel Software setups based on your system configuration.

6. Follow the instructions on the screen.

7. Set the host adapter physical address. The Fibre Channel AL Setup Utility will set up the required NT registry parameters and prompt you to select an ALPA for each of your host bus adapters. Use 01 for Host A and 02 for Host B for the host bus adapters.

8. Restart the system.

9. Repeat steps 1 through 8 on the other server.

Verifying Controller Properties

Perform the following procedures to verify controller properties.

Displaying SCSI Disks in the Devices Windowpane

To display SCSI disks in the devices windowpane:

1. Click **Start** on the taskbar.

2. Select **Programs, Command Console, CLI Window.**

3. Select **Serial** and click **OK.**

4. Select a baud rate of **9600,** and click **Connect.**

   NOTE: If a pop-up window reading "Error scanning subsystem" displays, wait two minutes and retry the connection.

   SWCC finds installed disks and displays them in a grid.

Verifying Properties

To verify the controller properties:

1. Double-click on a controller icon in the **Storage Window** to display **Controller Properties.**

2. Click on the tabs of the **Controller Properties** screen to confirm that the values in Table 2-4 are set.

Table 2-4: Controller Properties

Tab Name	Values
General	Allocation class: 0
	SCSI version: SCSI-2 or SCSI-3
Host Ports	Host Port 1 Requested Topology: Loop Hard
	Host Port 1 Actual Topology: Loop Hard
	Host Port 1 Requested Port Address: 71
	Host Port 1 Actual Port Address: 71
	Host Port 2 Requested Topology: Offline
	Host Port 2 Actual Topology: Offline
Cache	Cache flush time (seconds): 10
	Respond to internal cache battery condition: selected
Command Console LUN	SCSI-2 mode—Confirm that the screen is grayed out (disabled); if not, return to the "Disabling the Command Console LUN" section.
	SCSI-3 mode—Command Console LUN will be LUN 0 and cannot be disabled.
Connections	Operating System: WINNT
	Unit Offset: 0 for Port 1
Battery	Confirm that the battery is fully charged.

3. Click **OK** to close the **Controller Properties** window.

Configuring Large LUNs

The maximum logical units a host can access per controller port or controller port pair, if using redundant controllers, is 64. This access is accomplished automatically with ACS 8.6 and ACS 8.7. Older versions of ACS software required the StorageWorks Large LUN Utility to be run for a host to access more than eight LUNs per controller port or controller port pair if using redundant controllers. The Large LUN Utility is located in the Platform Kit that ships with the MA8000 product.

LUN 0 needs to be available and not reserved for Large LUN to function properly. Therefore it is best not to use LUN 0 for a cluster disk. SCSI-3 mode is preferred since LUN 0 is the Command Console LUN (CCL) and is not available.

Note that since mount points are not supported, cluster disks must have drive letters associated with them. Therefore a cluster is able to support a maximum of 22 cluster disks assuming drive letters A, B, C, and D are not used for cluster disks.

For maximum availability of LUNs, use the following configurations:

- Microsoft Windows NT Server 4.0, Enterprise Edition

  — SCSI-3 mode with CCL always enabled

  — LUNs beginning with D1 through D64

- Windows 2000 Advanced Server

  — SCSI-3 mode with CCL always enabled

  — LUNs beginning with D1 through D64

Configuring a Storageset

To configure a storageset:

NOTE: If errors are encountered while creating multiple storagesets, refer to the installation troubleshooting tips in Chapter 6.

1. Select **Storage** in the **Storage** window menu line.

2. Select **Add Virtual Disk** to begin step 1 of the Add Virtual Disk Wizard.

3. Select one of the available RAID level settings.

4. Click **Next** for step 2 of the Add Virtual Disk Wizard.

5. Select the devices you want to include in the virtual disk by clicking on the disks listed in the **Available Storage** window.

6. Click **Next** for step 3 of the Add Virtual Disk Wizard.

7. Use the displayed value for the virtual disk.

   IMPORTANT: All logical partitions based in a RAID set must be in the same cluster group.

8. Click **Next** for step 4 of the Add Virtual Disk Wizard.

9. Enter the virtual disk name.

10. Verify that the **Save controller configuration to virtual disk** check box is selected.

11. Click **Next** for step 5 of the Add Virtual Disk Wizard. Step 5 will display a summary of your selections.

    NOTE: If you are not satisfied with your selections, return to the applicable Wizard step using the **Back** button. When you are satisfied with your choices, click **Finish.**

12. Click **Finish** to create the virtual disk.

    The **Storage** window displays the virtual disk you created. The hourglass on the disk icon indicates that the storageset is being initialized. The drives you used to create the RAID set are highlighted in the **Devices** window.

13. Repeat steps 1 through 12 to create other virtual disks.

14. Wait for the virtual disk initialization to complete before proceeding (this may take some time, depending on the size and RAID level of the virtual disk created).

15. Close the **Storage** window.

Setting Up Connections

To set up HBA connections:

NOTE: Perform the following steps on one HBA at a time for an enhanced configuration.

1. Type the following and press **Enter:**

   ```
   show connections
   ```

   NOTE: Each connection must be deleted and reestablished in order to properly configure the HA/F500. Perform the following step if the connections were previously configured.

2. Type the following and press **Enter:**

   ```
   delete connection !NEWCONXX
   ```

   NOTE: The *XX* in !NEWCON*XX* refers to the connection number.

3. Repeat the above CLI command for each connection.

4. Close the CLI window.

5. Reboot the server you are currently working on.

6. Click **Start** on the Windows taskbar after rebooting.

7. Select **Programs, Command Console, CLI Window.**

8. Click **CLI Window** to display the **Connection Selection** dialog box.

9. Select **Serial** and click **OK** to display the **Connect Serial** dialog box.

10. Select the COM port (the port to which the storage controller is connected).

11. Select a baud rate of **9600.**

12. Click **Connect** to display the CLI window.

13. Press **Enter** to initialize communications.

14. Type the following command and press **Enter:**

    ```
    show connections
    ```

    NOTE: There should only be one connection, if you have a basic configuration. If not, delete all connections, reboot the server, and repeat the process.

15. Type the following command and press **Enter** to rename the single connection to make it available to the cluster.

    NOTE: The *XX* in !NEWCON*XX* refers to the connection number. Use a naming convention that will be easy to understand and remember. For example: NODE1T where Node1 is the name of the server. T stands for the top storage controller in the storage subsystem. T assumes that the maintenance cable is plugged in the storage controller that resides in the uppermost slot of the controller slots.

    ```
    rename !NEWCONXX NODE1T
    ```

16. Type the following command and press **Enter:**

    ```
    show connections
    ```

17. Verify that the new connection name is properly set.

18. Boot the second server. Log in to the domain when the login screen displays.

19. Go back to the CLI window, type the following command and press **Enter:**

    ```
    show connections
    ```

20. If the connection for the second server is not present, reboot the second server and try again.

21. Type the following command and press **Enter** to rename the connection for the second server.

   NOTE: The *XX* in !NEWCON*XX* refers to the connection number. Use a naming convention that will be easy to understand and remember. For example: NODE2T where Node2 is the name of the server. T stands for the top storage controller in the storage subsystem. T assumes that the maintenance cable is plugged in the storage controller that resides in the uppermost slot of the controller slots.

   ```
   rename !NEWCONXX NODE2T
   ```

22. Type the following and press **Enter:**

   ```
   show connections
   ```

23. Verify that the new connection name is properly set.

Verifying Access

To verify access to a storageset:

1. Be sure that the second server is powered down.

2. Click **Start** on the Windows taskbar.

3. Select **Programs, Command Console, CLI Window.**

4. Click **CLI Window** to display the **Connection Selection** dialog box.

5. Select **Serial** and click **OK** to display the **Connect Serial** dialog box.

6. Select the COM port (the port to which the storage controller is connected).

7. Select a baud rate of **9600.**

8. Click **Connect** to display the CLI window.

9. Press **Enter** to initialize communications.

10. Type the following and press **Enter:**

   ```
   show units full
   ```

11. Type the following command and press **Enter:**

    ```
    set unit name disable_access_path=ALL
    ```

    NOTE: The ALL access must be removed and set to NONE before adding the connection names defined earlier.

12. Type the following command and press **Enter:**

    ```
    show units full
    ```

13. Verify that the access is now set to **NONE.** If not, repeat the command.

    NOTE: Access will now consist of the connection names defined earlier.

14. Type the following command and press **Enter:**

    ```
    set unit name Enable_Access_Path=connection name
    ```

    Example: `set d0 Enable_Access_Path=NODE1T`

15. Repeat step 14 for the second server connection name:

    Example: `set d0 Enable_Access_Path=NODE2T`

16. Type the following command and press **Enter:**

    ```
    show units full
    ```

17. Verify that the connection names have been properly set. If not, repeat the commands.

18. Repeat steps 11 through 16 for each logical unit to be used in the cluster configuration.

19. Close the CLI window.

20. Reboot the server.

Volume Creation

To create storageset volumes:

1. Run Disk Administrator or Disk Management from one server. Make sure the second server is powered off.

   NOTE: To prevent Windows NT or Windows 2000 from reassigning the local server disks to different drive letters, make them "sticky" by assigning them the drive letters they currently have. This will also prevent drive C: from being renamed to a higher letter.

2. Create a volume on each storageset (virtual disk).

   NOTE: Create only one volume per storageset.

3. Format the volume with the NTFS file system.

4. Assign drive letters. Note the drive letter assignments for later use.

5. Close Disk Administrator or Disk Management.

6. Power off the first server.

7. Power on the second server.

8. Run Disk Administrator or Disk Management from the other server.

9. Assign the same drive letters to the storageset volumes that you assigned on the first server (step 4). The drive letters may not be applied until the server restarts.

   NOTE: The drive letter assignments for the storagesets must be the same on both servers.

10. Restart both servers.

11. After both servers restart, open Disk Administrator or Disk Management on both servers and verify that the drive letters are configured correctly.

Cluster Installation and Verification

The following sections install and verify the cluster installation.

Installing the Cluster Software

Install and configure Microsoft cluster software on the host servers as described in the manual provided with Microsoft Windows NT Server 4.0, Enterprise Edition, Windows 2000 Advanced Server, or Windows 2000 Datacenter Server.

Verifying Creation of the Cluster

To verify the creation of the cluster:

1. Type the following CLI command and press **Enter** to shut down your storage subsystem.

   ```
   shutdown this_controller
   ```

2. Shut down and power off both servers.

3. Power off the storage subsystem and then power it back on.

4. Power both servers back on.

5. From the Windows desktop on either clustered server, select **Start, Programs, Administrative Tools (Common), Cluster Administrator.**

6. Enter the name or IP address of the cluster when you are prompted for **Cluster Name.** If you run Cluster Administrator on a cluster node, enter a period and the cluster is found automatically.

   If the cluster has been created correctly, the computer names of both cluster nodes appear on the left side of the **Cluster Administrator** window.

7. If the cluster is not working correctly, refer to the installation troubleshooting tips in Chapter 6.

Verifying Node Failover

> **IMPORTANT:** Do not run any client activity while testing failover events.

To verify failover of a cluster node:

1. From the desktop on both servers, select **Start, Programs, Administrative Tools (Common), Cluster Administrator.**

2. When you are prompted for the **Cluster Name** or **Server Name,** enter the name or IP address of the cluster.

3. Be sure that all cluster resources and cluster groups are online. Verify that some of the cluster groups are owned by the server you will be powering off so that a failure event will result in a failover of cluster groups.

4. Power off the server mentioned in step 3.

   Within several seconds, Cluster Administrator running on the surviving node should bring online all the predefined resources and groups that were previously owned by the powered-off server. If, after a minute, nothing appears to have occurred, refresh the screen by selecting **Refresh** (or by pressing the **F5** key).

5. If failover is not working correctly, refer to the installation troubleshooting tips in Chapter 6.

6. Power on the powered-off server to continue.

Verifying Network Client Failover

Now that you have verified that the server is correctly running as a cluster node, the next step is to verify that network clients can interact with the cluster.

To verify network client failover:

1. Be sure that both cluster nodes are running and verify through Cluster Administrator that all groups and resources are online.

2. For each hard disk in the shared storage, Microsoft cluster software automatically creates a cluster group that consists of a single resource, the disk drive. Using Cluster Administrator, add an unused IP address as another resource to one of these groups. (Do **not** use the Cluster Group.) Bring the newly created IP resource online.

3. Open a Command Prompt window on a network client machine.

4. Be sure that the network client can access the IP address resource. Regardless of whether you are using WINS or DHCP, you can execute the PING command to check the connection.

   From the network client, execute a PING command using the cluster IP address as the argument. The client has successfully accessed the IP address resource if you get a response similar to:

   ```
   Reply from IP Address: bytes=xx time=xxxms TTL=xx
   ```

   The client has not successfully accessed the cluster resource if you get a response of:

   ```
   Reply from IP Address: Destination host unreachable
   ```

5. Following the successful completion of the PING command, use Cluster Administrator to perform a manual failover of the cluster group that contains the IP address resource.

6. After the manual failover completes, execute the PING command again.

   As soon as the other node brings the cluster group online, a response similar to the one noted in step 4 should be returned. If the client successfully accessed the failed-over IP address resource, the cluster is working. If the client was unsuccessful, either the cluster group was not configured correctly, the failover did not occur, or the PING command was performed before the failover activity completed.

7. If network client failover is not working correctly, refer to the installation troubleshooting tips in Chapter 6.

8. If you want to verify a more extreme case, rather than failover the IP address resource, power off the primary cluster node and verify that the resource fails over to the other node.

Setting Up Cluster Groups and Cluster Resources

The clustering hardware is now set up and configured. The Microsoft operating system with Microsoft cluster software is installed on both servers. You have done minimal validation of the cluster and everything is working as planned. Now it is time to set up the applications to use the capabilities of clustering.

If you have defined your cluster needs correctly and determined how to fit the applications and environment into a cluster, then configuring the applications will be a straightforward task.

Although details of these procedures are beyond the scope of this guide, several documents are available from both Microsoft and Compaq to assist you with these final steps.

The best source of information concerning the steps to configure cluster groups and cluster resources is the *Microsoft Cluster Server Administrator Guide,* which provides detailed setup information.

Compaq has worked directly with several application vendors throughout the development of ProLiant Clusters. As a result of these efforts, Compaq has provided a number of integration documents to assist you with installing these applications in a ProLiant Cluster environment.

Integration documents are available at www.compaq.com/highavailability.

Setting Up Dual Storage Controllers

This section describes how to convert a system from a single-adapter, single-storage controller configuration to a single-adapter, dual-storage controller configuration. Dual controller configuration allows failover to occur between the controllers in the event of a controller failure.

Hardware Setup

To set up the hardware:

1. Click **Start** on the taskbar.

2. Select **Programs, Command Console, CLI Window.**

3. Click **CLI Window** to display the **Connection Selection** dialog box.

4. Select **Serial** and click **OK** to display the **Connect Serial** dialog box.

5. Select the COM port (the port to which the storage controller is connected).

6. Select a baud rate of **9600.**

7. Click **Connect** to display the CLI window.

8. Press **Enter** to initialize communications.

9. Type the following command and press **Enter:**

   shutdown this_controller

10. Close the CLI window.

11. Shut down and power off both host servers.

12. Power off the storage subsystem.

13. Install the second storage controller into the empty slot.

Figure 2-4: Installing the second storage controller (RA8000 shown)

14. Install the PCMCIA card in the controller.
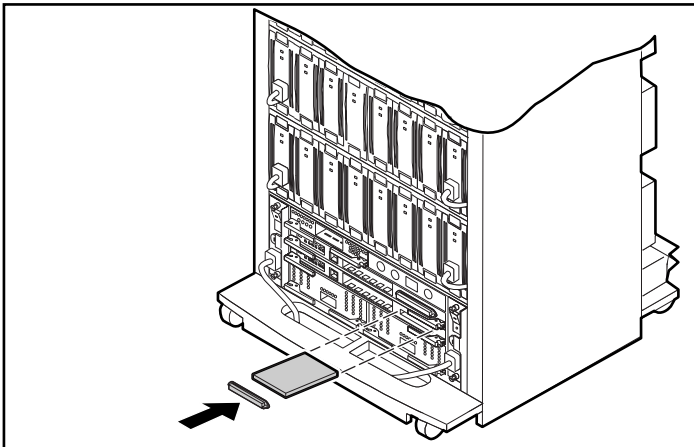
Figure 2-5: Installing the PCMCIA card into the second storage controller (RA8000 shown)

15. Connect port 1 (the left port) of the second controller to the hub.

16. Power on the storage subsystem.

17. Power on both host servers.

Configuring Controllers for Failover Mode

NOTE: To use two storage controllers and one HBA in each server, you need to configure the controllers to Transparent Failover mode.

1. Click **Start** on the taskbar.

2. Select **Programs, Command Console, CLI Window.**

3. Click **CLI Window** to display the **Connection Selection** dialog box.

4. Select **Serial** and click **OK** to display the **Connect Serial** dialog box.

5. Select the COM port (the port to which the storage controller is connected).

6. Select a baud rate of **9600.**

7. Click **Connect** to display the CLI window.

8. Press **Enter** to initialize communications.

9. Type the following command and press **Enter:**

   ```
   set failover copy=this_controller
   ```

10. Wait two minutes for the controllers to reset.

11. Verify that both controllers are in dual redundant configuration (failover) mode by entering the following commands and viewing the results:

    ```
    show this_controller
    ```

    ```
    show other_controller
    ```

12. Close the CLI window.

Verifying the System

Restart both servers and inspect the cluster and the event logs of the hosts to verify that the system is functioning without errors.

Installing Additional Storage Subsystems

The following sections provide information on installing an additional storage subsystem to the cluster.

Setting Up a Single Cluster with an Additional Storage Subsystem (Single Controller Configuration)

1. Set up the cluster with the first storage subsystem according to the instructions starting on page 2-4.

2. To add a second storage subsystem, follow the installation instructions starting on page 2-4 with the **important exceptions** included in Table 2-5.

Table 2-5: Setting Up a Single Cluster with an Additional Storage Subsystem (Single Controller Configuration)

Chapter 2 Section Heading	Page	Action
Setting Up the Servers	2-3	Omit
Installing the Host Bus Adapter	2-5	Omit
Installing the StorageWorks Command Console Client	2-8	Omit
Installing the Adapter Driver	2-11	Omit
Running the FC-AL Setup Utility	2-12	Omit
Cluster Installation and Verification	2-22	Omit this and all following sections.

NOTE: Refer to Chapter 5 for information on adding shared storage.

Setting Up a Single Cluster with an Additional Storage Subsystem (Dual Controller Configuration, Transparent Failover or Multibus Failover)

1. Set up the cluster with the first storage subsystem according to instructions starting on page 2-4.

2. To add a second storage subsystem, follow the installation instructions starting on page 2-4 with the **important exceptions** included in Table 2-6.

Table 2-6: Setting Up a Single Cluster with an Additional Storage Subsystem (Dual Controller, Transparent Failover or Multibus Failover)

Chapter 2 Section Heading	Page	Action
Setting Up the Servers	2-3	Omit
Installing the Host Bus Adapter	2-5	Omit
Installing the StorageWorks Command Console Client	2-8	Omit
Installing the Adapter Drivers	2-11	Omit
Running the FC-AL Setup Utility	2-12	Omit
Cluster Installation and Verification	2-22	Omit
Setting Up Dual Storage Controllers	2-26	**Execute this step.**

NOTE: Refer to Chapter 5 for information on adding shared storage.

Multiple Cluster Configurations for Hubs

This section contains instructions for setting up two clusters with hubs on the same storage system. This section will refer to steps that can be found in Chapters 2 and 4 of this guide.

Key Considerations

When setting up multiple cluster configurations with hubs, be aware of the following key considerations:

- Virtual disk naming conventions

- ALPA assignment

- Server-to-storage connection paths

Virtual Disk Naming Conventions

Table 2-7 shows the naming conventions used with the Large LUN Utility.

Table 2-7: Storage Controller Configurations

	Unit Offset	Drive Labels
Cluster 1	0	D1-D64
Cluster 2	100	D101-D164

ALPA Assignment

All HA/F500 configurations require a LOOP_HARD topology. Changing the port topology to LOOP_HARD requires an ALPA for the controller port. The ALPA is used to identify the device on the arbitrated loop whenever the loop initialization process (LIP) occurs. Having an ALPA value assigned increases the likelihood that the value will be available for that device when the LIP occurs. If a value is not chosen or is not available, the initialization process searches through the entire table of acceptable ALPA values until it finds a value that is available for use by the device.

Server-to-Storage Connection Paths

Server-to-storage connection paths must be defined when configuring multiple cluster configurations for multibus failover. If connections are not defined, all new connections (the automatically generated "!NEWCON*XX*" connections) will be assigned Unit Offset 0, and the cluster using port 2 of each controller will not be able to access the logical units intended for its use. The connection path setting helps segregate the logical units between the two clusters.

NOTE: The *XX* in !NEWCON*XX* refers to the connection number.

Assigning server-to-storage connection paths is necessary only in the multibus failover configuration (Figure 2-6).

Multiple Cluster Configurations

The following illustrations depict two multiple cluster configurations for use with a storage subsystem.

Figure 2-6 shows two basic cluster configurations connected to a single storage controller in a single storage subsystem. The storage controller is configured as active/active.
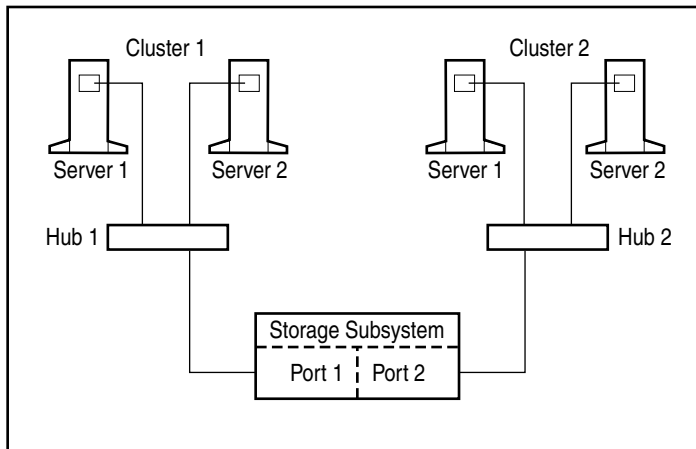
Figure 2-6: Multiple cluster configuration (single storage controller)

Figure 2-7 shows two enhanced cluster configurations connected to dual storage controllers in a single storage subsystem. The storage controllers are configured as active/active (multibus failover).
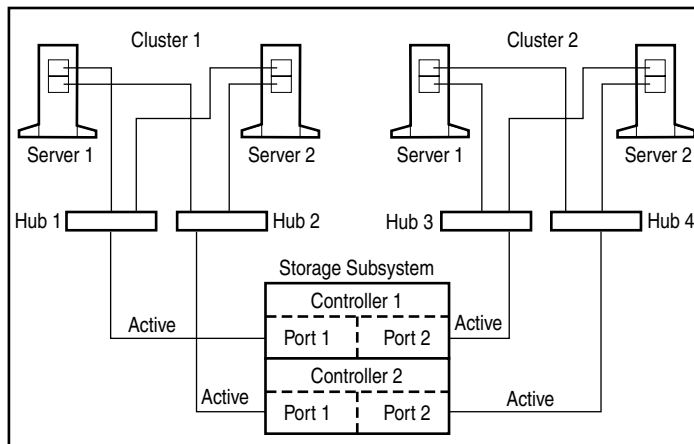
Figure 2-7: Multiple cluster configuration (multibus failover)

Setting Up Single Storage Controller Configurations

To set up a multiple cluster configuration with a single storage controller (refer to Figure 2-6):

1. Set up the first cluster according to instructions starting on page 2-3.

2. After you set up the first cluster, follow the steps on page 2-10 of this chapter to set the port topology and the ALPA settings for the storage controller.

3. Follow the cluster installation instructions in this chapter to set up the second cluster, with the **important exceptions** included in Table 2-8.

Table 2-8: Setting Up the Second Cluster (Single Controller Configuration)

Chapter 2 Section Heading	Page	Action
Installing the Disk Drives, Storage Controller, and Program Card	2-4	Omit
Set Controller to No Failover	2-8	Omit
Disabling the Command Console LUN	2-9	Omit
Setting Port Topologies and the Arbitrated Loop Physical Address	2-39	Follow the instructions on page 2-39 of this chapter instead of those on page 2-10.
Discovering Devices	2-11	Omit
Running the FC-AL Setup Utility	2-12	In step 7, use 04 for Host A and 08 for Host B.
Verifying Controller Properties	2-13	In Table 2-4, use the following text for Host Ports: • Host Port 2 Requested Topology: Loop Hard • Host Port 2 Requested Port Address: 72 • Host Port 2 Actual Port Address: 72
Configuring a Storageset	2-16	Use the following text for step 9: Enter the virtual disk name. Windows NT recognizes D100 through D107 for the second cluster if the Large LUN Utility is not installed.
Setting up Dual Storage Controllers	2-26	Omit

Setting Up Dual Redundant Controllers (Transparent Failover Configuration)

To set up a multiple cluster configuration with dual redundant controllers (transparent failover):

1. Set up the first cluster according to instructions starting on page 2-3.

2. After you set up the first cluster, follow the steps on page 2-10 of this chapter to set the port topology and the ALPA settings for the storage controller.

3. Follow the cluster installation instructions in this chapter to set up the second cluster, with the **important exceptions** included in Table 2-9.

Table 2-9: Setting Up the Second Cluster (Transparent Failover Configuration)

Chapter 2 Section Heading	Page	Action
Installing the Disk Drives, Storage Controller, and Program Card	2-4	Omit
Set Controller to No Failover	2-8	Omit
Disabling the Command Console LUN	2-9	Omit
Setting Port Topologies and the Arbitrated Loop Physical Address	2-39	Follow the instructions on page 2-39 of this chapter instead of those on page 2-10.
Discovering Devices	2-11	Omit
Running the FC-AL Setup Utility	2-12	In step 7, use 04 for Host A and 08 for Host B.
Verifying Controller Properties	2-13	In Table 2-4, use the following text for Host Ports: • Host Port 2 Requested Topology: Loop Hard • Host Port 2 Requested Port Address: 72 • Host Port 2 Actual Port Address: 72
Configuring a Storageset	2-16	Use the following text for step 9: Enter the virtual disk name. Windows NT recognizes D100 through D107 for the second cluster if the Large LUN Utility is not installed.

Setting Up Dual Redundant Controllers (Multibus Failover)

To set up a multiple cluster configuration with dual redundant controllers (multibus failover), refer to Figure 2-7 for an example of this configuration. Since all four ports are active, the unit offsets are not assigned automatically and must be configured manually.

1. Set up the first cluster according to instructions in Chapters 2 and 4.

2. After you set up the first cluster, follow the steps on page 2-10 of this chapter to set the port topology and the ALPA settings for the storage controllers.

3. Follow the cluster installation instructions in Chapters 2 and 4 to set up the second cluster, with the **important exceptions** included in Table 2-10.

Table 2-10: Setting Up the Second Cluster (Multibus Failover Configuration)

Chapters 2 and 4 Section Heading	Page	Action
Installing the Disk Drives, Storage Controller, and Program Card	2-4	Omit
Set Controller to No Failover	2-8	Omit
Disabling the Command Console LUN	2-9	Omit
Setting Port Topologies and the Arbitrated Loop Physical Address	2-39	Follow the instructions on page 2-39 of this chapter instead of those on page 2-10
Discovering Devices	2-11	Omit
Running the FC-AL Setup Utility	2-12	In step 7, use 04 for Host A and 08 for Host B.

continued

Table 2-10: Setting Up the Second Cluster (Multibus Failover Configuration)
continued

Chapters 2 and 4 Section Heading	Page	Action
Verifying Controller Properties	2-13	In Table 2-4, use the following text for Host Ports:
		• Host Port 2 Requested Topology: Loop Hard
		• Host Port 2 Requested Port Address: 72
		• Host Port 2 Actual Port Address: 72
Configuring a Storageset	2-16	Use the following text for step 9: Enter the virtual disk name. Windows NT recognizes D100 through D107 for the second cluster if Large LUN Utility is not installed.
Volume Creation	2-21	Follow the instructions in this section to assign the connection paths for the **primary** path of the second cluster (page 2-40). These instructions replace the Volume Creation steps on page 2-21.
Establishing a Redundant Path for the Microsoft Cluster	4-8	Step 2: Omit
		Step 5: Disregard the Note section and use port 2 (the right port) of each controller.
		After completing step 6, follow the instructions in this section to assign connection paths for the redundant path of the second cluster (page 2-41).

Referenced Procedures

The following three procedures are referenced in this chapter as part of the configuration setup instructions. Apply these procedures only as indicated by the instructions and tables in this chapter.

Setting Port Topology and the Arbitrated Loop Physical Address

NOTE: Please note that you are now setting port topology to port 2.

1. From the CLI window, type the following command and press **Enter:**

   ```
   set this_controller port_2_al_pa=72
   ```

2. Type the following and press **Enter:**

   ```
   set this_controller port_2_topology=loop_hard
   ```

   IMPORTANT: Omit step 3 when setting up a single controller configuration.

3. Type the following and press **Enter:**

   ```
   restart other_controller
   ```

4. Type the following and press **Enter:**

   ```
   restart this_controller
   ```

5. Wait two minutes for the controllers to restart.

6. Type the following and press **Enter:**

   ```
   show this_controller
   ```

   NOTE: This command assumes that the maintenance port cable is connected to the top controller of the pair.

7. Verify the following:

   ```
   port_2_topology=loop_hard (standby)
   port_2_al_pa=72 (72 negotiated)
   ```

   IMPORTANT: Omit steps 8 and 9 when setting up a single controller configuration.

8. Type the following and press **Enter:**

   ```
   show other_controller
   ```

   NOTE: This command assumes that the maintenance port cable is connected to the top controller of the pair.

9. Verify the following:

   ```
   port_2_topology=loop_hard (loop up)
   ```

   ```
   port_2_al_pa=72 (72 negotiated)
   ```

   NOTE: Refer to the HSG80 documentation if you need more information about the port topology and ALPA settings.

Assigning Connection Paths for the Primary Path of the Second Cluster

IMPORTANT: If connections are not defined, all new connections (the automatically generated "*!NEWCONXX*" connections) will be assigned Unit Offset 0. This will result in all boxes seeing all drives, minus units that have been reserved by the other cluster. The unit will be offline in the Disk Administrator.

NOTE: Only the connections in the cluster using port 2 of each controller need to be renamed. The cluster using port 1 of each controller does not require renaming.

1. Power off **all** servers connected to the disk array.

2. Power up Server 1 in Cluster 2. (This server should have ALPA 4.)

3. Power up Server 2 in Cluster 2. (This server should have ALPA 8.)

   NOTE: Do **not** power on any server in Cluster 1.

4. From the CLI window, type the following command and press **Enter:**

   ```
   show connections
   ```

5. Write down the connections that are OL THIS or OL OTHER (OL = Online) for renaming the connection in step 7.

6. Check the ALPA value of each connection. It should be 4 for Server 1 and 8 for Server 2 from Cluster 2. If other connections are online, ignore them. The ALPA is displayed as a six-digit number in the ADDRESS column (000004 or 000008).

7. Rename each connection by typing the following command and pressing **Enter:**

   ```
   rename !newconxx xxxxxxxxx
   ```

   NOTE: You may use a maximum of nine alphanumeric characters.

   ```
   Example: rename !NEWCONXX CLU2NODE1
   ```

8. Type the following command and press **Enter:**

   ```
   set unit name disable_access_path=ALL
   ```

   NOTE: The ALL access must be removed and set to NONE before adding the connection names defined earlier.

9. Set the unit offset on each connection by typing the following command and pressing **Enter:**

   ```
   set xxxxxxx unit_offset=100
   ```

10. Type the following and press **Enter:**

    ```
    show connections
    ```

11. Verify the names and unit offsets assigned previously.

When step 10 has been completed, return to Table 2-10 of this chapter and continue with the instructions for "Establishing a Redundant Path for the Microsoft Cluster" in Chapter 4 of this guide.

Assigning Connection Paths for the Redundant Path of the Second Cluster

IMPORTANT: If connections are not defined, all new connections (the automatically generated "*!NEWCONxx*" connections) will be assigned Unit Offset 0. This will result in all boxes seeing all drives, minus units that have been reserved by the other cluster. (The unit will be offline in Disk Administrator.)

NOTE: Only the connections in the cluster using port 2 of each controller need to be renamed. The cluster using port 1 of each controller does not require renaming.

1. Power off **all** servers connected to the disk array.

2. Power on Server 1 in Cluster 2. (This server should have ALPA 4.)

3. Power on Server 2 in Cluster 2. (This server should have ALPA 8.)

   NOTE: Do **not** power on any server in Cluster 1.

4. From the CLI window, type the following command and press **Enter:**

   ```
   show connections
   ```

5. Write down the connections that are OL THIS or OL OTHER (OL = Online) for renaming the connection in step 7.

6. Check the ALPA value of each connection. It should be 4 for Server 1 and 8 for Server 2 from Cluster 2. If other connections are online, ignore them. The ALPA is displayed as a six-digit number in the ADDRESS column (000004 or 000008).

7. Rename each connection by typing the following command and pressing **Enter:**

   ```
   rename !newconxx xxxxxxxxx
   ```

   NOTE: You may use a maximum of nine alphanumeric characters.

   Example: ```rename !newconxx CLU2NODE1```

8. Set the unit offset on each connection by typing the following command and pressing **Enter:**

   ```
   set xxxxxxxx unit_offset=108
   ```

9. Type the following and press **Enter:**

   ```
   show connections
   ```

10. Verify the names and unit offsets assigned previously.

3

Setting Up the ProLiant Cluster HA/F500 with Switches

Preinstallation Instructions

Before setting up the ProLiant Cluster HA/F500 with switches, verify that the hardware and software kits are appropriate for this installation. The following section gives more details about the different installation configurations.

Checking the Hardware and Software

Verify that the appropriate hardware and software installation kits were received.

In the following tables, the numbers in parentheses refer to configurations with *N* (where *N* = 2, 3, or 4) storage subsystems.

Table 3-1: Basic Configuration with Switches (Single Storage Controller)

Quantity	Description
2	ProLiant servers
2	Host bus adapters
1(*N*)	Storage subsystem with one storage controller per subsystem
1	Storage controller platform kit
1(*N*)	ACS Controller software
1	8- or 16-port Fibre Channel switch
1	ProLiant Cluster HA/F500 Basic Kit with Ethernet crossover cable for the server interconnect
1	Microsoft operating system, which includes Microsoft cluster software
As required	Optical Fibre Channel interconnect hardware

Table 3-2: Basic Configuration with Switches (Dual Storage Controllers)

Quantity	Description
2	ProLiant servers
2	Host bus adapters
1(*N*)	Storage subsystem with two storage controllers per subsystem
1	Storage controller platform kit
2(2*N*)	ACS Controller software
1	8- or 16-port Fibre Channel switch
1	ProLiant Cluster HA/F500 Basic Kit with Ethernet crossover cable for the server interconnect
1	Microsoft operating system, which includes Microsoft cluster software
As required	Optical Fibre Channel interconnect hardware

Table 3-3: Enhanced Configuration for Switches

Quantity	Description
2	ProLiant servers
4	Host bus adapters
1(*N*)	Storage subsystem with two storage controllers per subsystem
1	Storage controller platform kit
2(2*N*)	ACS Controller software
2	8- or 16-port Fibre Channel switches
1	ProLiant Cluster HA/F500 Enhanced Kit containing Ethernet crossover cable for the server interconnect and Secure Path Software
1	Microsoft operating system, which includes Microsoft cluster software
As required	Optical Fibre Channel interconnect hardware

Table 3-4: Datacenter Configuration with Switches

Quantity	Description
4	ProLiant servers
8	Host bus adapters
1	Storage subsystem with two storage controllers per subsystem
1	Storage controller platform kit
2	ACS Controller software
2	8- or 16-port Fibre Channel switches
1	ProLiant Cluster HA/F500 Enhanced Kit, includes two Secure Path Software licenses
4	Microsoft operating system, which includes Microsoft cluster software
As required	Optical Fibre Channel interconnect hardware
As required	Additional Secure Path software packages, one additional Secure Path license per node

If you are missing any required component, contact your local sales representative or call 1-800-OK-COMPAQ.

Installing the Hardware

The following sections provide information on installing the hardware components used in a ProLiant cluster.

Setting Up the Servers

Follow the installation instructions in the Compaq ProLiant server documentation to set up each server as a stand-alone server. Then use the instructions in the following sections to configure the cluster.

Setting Up the Storage Subsystem

Refer to the documentation that was shipped with the storage subsystem for detailed installation instructions.

Installing the Disk Drives, Storage Controllers, and Program Cards

The ProLiant Cluster HA/F500 uses hard drives, storage controllers, and PCMCIA program cards that must be installed in the storage subsystem (see Figure 3-1).

IMPORTANT: If you are setting up a basic configuration, verify that only one controller is installed. If more than one controller is present, remove the additional controllers to facilitate installation.

To install a disk drive (1):

1. Insert a disk drive into the shelf guide slot.

2. Slide the disk drive into the shelf and use the attached lever to push it into place.

   NOTE: For optimum SCSI bus distribution, install the disk drives along different SCSI buses.

To install a storage controller (2):

1. Insert the storage controller into the top shelf guide slot.

2. Slide the storage controller into the shelf until the mounting tabs snap into place.

To install a PCMCIA program card (3) in the storage controller:

1. Remove the program card cover (4) from the controller PCMCIA slot.

2. Insert the PCMCIA program card (3) into the top controller slot.

3. Replace the program card cover (4) over the controller slot.
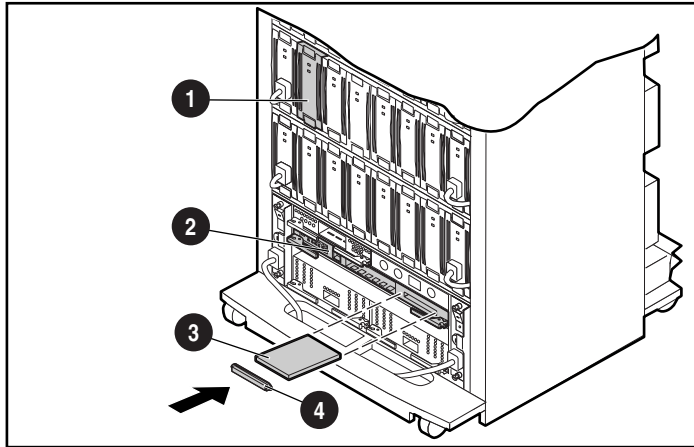
Figure 3-1: Installing the disk drives, storage controller, and PCMCIA Cards (RA8000 shown)

Installing the Host Bus Adapter

Follow the installation instructions in the server documentation to install the host bus adapter in the servers. Install one adapter in each server.

For the basic configurations (single host bus adapter, single or dual storage controllers), only one adapter is installed per server. Chapter 4 discusses the enhanced configuration (dual host bus adapters, dual storage controllers) in which a second adapter is installed.

Connect the adapter from each server to the Fibre Channel switch. Connect port 1 (the left port) of the storage controller to the Fibre Channel switch.

Designating the Server as a Maintenance Terminal

A server must be connected to the storage controller to provide a maintenance terminal (see Figure 3-2).

NOTE: Only one server should be designated as the maintenance terminal. It is recommended that a separate stand-alone server that is not part of the cluster be designated as the maintenance server.

To connect the server to the storage controller:

1. Connect the RJ-12 connector on the communications cable to the maintenance port (1) on the storage controller.

2. Connect the 9-pin serial connector on the communications cable to either the COM1 or COM2 port (2) on the server.

   NOTE: Note which serial port is used. This information will be needed when setting up the communications program and configuring the controller.
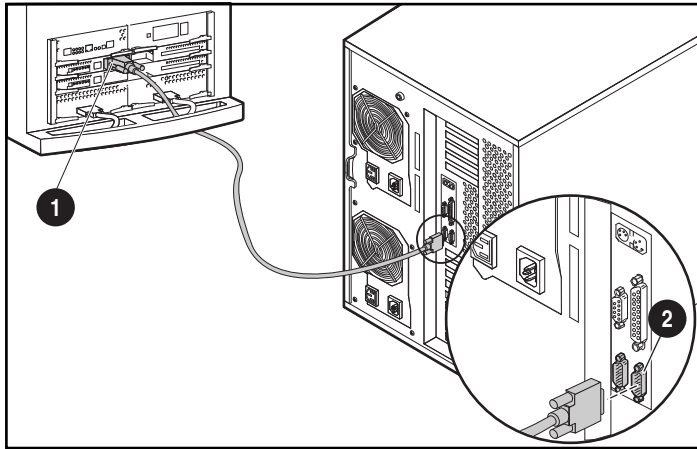
Figure 3-2: Designating a maintenance terminal

Turning On the Storage Subsystem Power

To power on the subsystem:

1. Connect the storage subsystem cabinet to an AC power outlet (1).

2. Turn the storage subsystem power to the on position (2). (This refers to RA8000/ESA12000 storage subsystems.)

3. Wait until the storage subsystem is completely booted and ready to operate.

4. Turn on both servers.

The storage subsystem is ready to operate when the Reset LED on the storage controller blinks at a rate of one time per second.
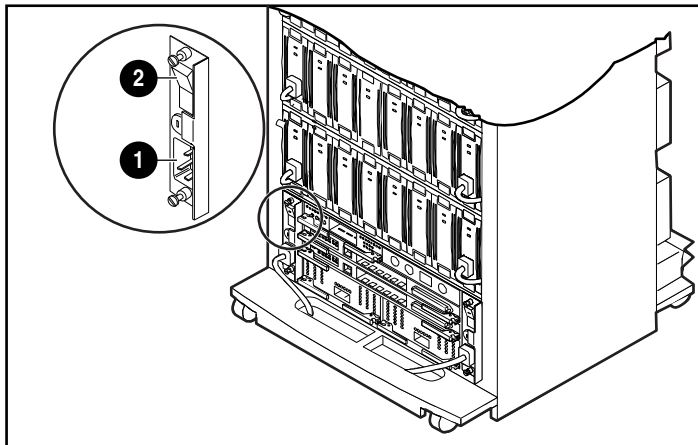
Figure 3-3: Turning on storage subsystem power (RA8000/ESA12000)

Setting Up the Switch

Refer to the following sections to set up the switch. For more information about the switch, refer to the documentation that shipped with the switch.

Installing the Switch

For current installation and configuration instructions, refer to the documentation included with the switch.

1. Connect the AC power cable to the switch AC connector.

2. Turn on the AC power switch. The Power-On Self-Test (POST) runs automatically.

   NOTE: Refer to the switch documentation for descriptions of POST encounter errors and LED indicators.

When each test completes successfully, the port LEDs display a steady green light.

For custom configuration information, refer to the documentation that comes with the switch. Continue with the following steps to complete the cluster installation.

Updating Switch Firmware

If you are using the StorageWorks Fibre Channel Storage Switches (8 or 16) in a ProLiant Cluster HA/F500 configuration, be sure you have configured the switches using the correct firmware version. Firmware updates can be downloaded from www.compaq.com/storage.

Combining Switch Models in a Cluster Configuration

Different SAN switch models can be combined in a cluster configuration. However, if you experience difficulty establishing communication between different switch models, perform the following steps:

1. Disable the SAN switches.

2. On the SAN switches, set the following parameter:

   vcencode mode = 1

3. Re-enable the SAN switches.

Installing the StorageWorks Command Console Client

Install the StorageWorks Command Console Client (SWCC) software on the host designated as the maintenance server.

To install the SWCC Client:

1. Insert the storage subsystem CD-ROM from the HSG80 NT Intel platform kit in the CD-ROM drive.

2. Select **Solution Software for ACS.**

3. Click **Skip Driver Upgrade.**

4. Select **StorageWorks Command Console (SWCC).**

5. Select **CLI, Command Window, HSG80 or Newer.**

6. Follow the instructions on the screen.

Configuring the Storage Subsystem for Fibre Channel

The storage subsystem is preconfigured with specific parameters. These parameters may need to be changed depending on which cluster configuration is chosen. This section will review the required steps to make these changes.

Setting Controller to No Failover

To set the storage controller to no failover:

1. Click **Start** on the taskbar.

2. Select **Programs, Command Console**, **CLI Window** to display the **Connection Selection** dialog box.

3. Select **Serial** and click **OK** to display the **Connect Serial** dialog box.

4. Select the COM port (the port to which the storage controller is connected).

5. Select a baud rate of **9600.**

6. Click **Connect** to display the CLI window.

    NOTE: You may need to press **Enter** to initialize communications before entering CLI commands.

7. Type the following command and press **Enter:**

    `show this_controller`

8. Verify that the "Not Configured for Dual-redundancy" text is displayed in the controller properties.

9. If not, type the following command and press **Enter:**

    `set nofailover`

    NOTE: If you are unable to connect with the storage controller, reset the controller by pressing the flashing Reset LED button on the controller, and retry the connection after two minutes.

Disabling the Command Console LUN

NOTE: The Command Console LUN is disabled by default for SCSI mode SCSI-2. For SCSI mode SCSI-3 the Command Console LUN cannot be disabled.

To disable the Command Console LUN:

1. Click **Start** on the Windows taskbar.

2. Select **Programs, Command Console, CLI Window** to display the **Connection Selection** dialog box.

3. Select **Serial** and click **OK** to display the **Connect Serial** dialog box.

4. Select the COM port (the port to which the storage controller is connected).

5. Select a baud rate of **9600.**

6. Click **Connect** to display the CLI window.

7. Type the following command and press **Enter:**

   ```
   set this_controller nocommand_console_lun
   ```

8. After the prompt returns, type the following and press **Enter:**

   ```
   show this_controller
   ```

9. Verify that the "Command Console LUN disabled" text appears in the controller properties. If so, then continue with the next section.

Setting Port Topology

Depending on the configuration, you may need to change the port topology settings. To set port topology settings:

NOTE: Refer to the HSG80 documentation for more details on changing the port topology settings.

1. From the CLI window, type the following and press **Enter:**

   ```
   set this_controller port_1_topology=fabric
   ```

2. Type the following and press **Enter:**

   ```
   set this_controller port_2_topology=offline
   ```

3. Type the following and press **Enter:**

   ```
   show this_controller
   ```

4. Verify the following:

   ```
   port_1_topology=FABRIC (Fabric Up)
   port_2_topology=OFFLINE (Offline)
   ```

Discovering Devices

To enable the controller to discover the disk drives:

1. From the CLI window, type the following and press **Enter:**

   ```
   run config
   ```

2. Wait a few minutes for the controller to scan and discover the disk drives.

3. Type the following and press **Enter:**

   ```
   show devices
   ```

4. Verify that the devices discovered are consistent with the disk drives installed.

   The disk drives will be named in the DISK##### format. Refer to the storage subsystem documentation for the naming conventions for storage subsystem devices.

5. Close the CLI window.

Installing the Adapter Driver

To install the adapter driver:

IMPORTANT: If the host bus adapter drivers are being installed on a new system with a Windows NT operating system, you need to make sure an older version of the driver is installed on the system first to prevent an error with the upgrade utility. An older version of the host bus adapter driver can be downloaded from www.compaq.com/products/storageworks/adapters.html.
A new system with a Windows 2000 operating system does not require a previous driver version.

1. Insert the storage subsystem CD-ROM from the HSG80 NT Intel platform kit in the CD-ROM drive.

2. Select **Solution Software for ACS.**

3. Select **KGPSA Driver Upgrade.**

4. Follow the instructions on the screen.

5. Repeat steps 1 through 4 on the other servers.

6. Verify the HBA driver version installed on the server.

7. Go to www.compaq.com/products/storageworks/adapters.html and check for the latest Fibre Channel host bus adapter driver version supported.

   IMPORTANT: Skip the following step if you have the latest Fibre Channel HBA driver installed on the server.

8. Download the latest Fibre Channel HBA driver utility and install the utility on each server, one server at a time.

Running the Fibre Channel Fabric Setup Utility

IMPORTANT: Run this utility before any NTFS volumes are created on the storage subsystem. If NTFS volumes exist, running Fibre Channel Fabric Setup and rebooting may reassign drive letters. Refer to the storage subsystem documentation for details on drive letter reassignment.

The Fibre Channel Fabric Setup Utility installs the version of the HSZDISK driver required for Windows NT Clusters in a switched fabric environment.

To run the Fibre Channel Fabric Setup Utility:

1. Insert the storage subsystem CD-ROM from the HSG80 NT Intel platform kit in the CD-ROM drive.

2. Select **Solution Software for ACS.**

3. Click **Skip Driver Upgrade.**

4. Select **Fibre Channel Software Setup.**

5. Select one of the Fibre Channel Software setups based on your system configuration.

6. Follow the instructions on the screen.

7. Restart the system.

8. Repeat steps 1 through 7 on the other servers.

Verifying Controller Properties

Perform the following procedures to verify controller properties.

Displaying SCSI Disks in the Devices Windowpane

To display SCSI disks in the devices windowpane:

1. Click **Start** on the taskbar.

2. Select **Programs, Command Console, CLI Window.**

3. Select **Serial** and click **OK.**

4. Select a baud rate of **9600,** and click **Connect.**

   NOTE: If a pop-up window reading "Error scanning subsystem" displays, wait
   two minutes and retry the connection.

SWCC finds installed disks and displays them in a grid.

Verifying Properties

To verify the controller properties:

1. Double-click a controller icon in the **Storage** window to display **Controller Properties.**

2. Click the tabs of the **Controller Properties** screen to confirm that the values in Table 3-5 are set.

Table 3-5: Controller Properties for Switches

Tab Name	Values
General	Allocation class: 0
	SCSI version: SCSI-2 or SCSI-3
Host Ports	Host Port 1 Requested Topology: Fabric
	Host Port 1 Actual Topology: Fabric Up
	Host Port 2 Requested Topology: Offline
	Host Port 2 Actual Topology: Offline
Cache	Cache flush time (seconds): 10
	Respond to internal cache battery condition: selected
Command Console LUN	SCSI-2 mode—Confirm that the screen is grayed out (disabled); if not, return to the "Disabling Command Console LUN" section.
	SCSI-3 mode—Command Console LUN will be LUN 0 and cannot be disabled.
Connections	Operating System: WINNT
	Unit Offset: 0 for Port 1
Battery	Confirm that the battery is fully charged.

3. Click **OK** to close the **Controller Properties** window.

Configuring Large LUNs

The maximum logical units a host can access per controller port or controller port pair, if using redundant controllers, is 64. This access is accomplished automatically with ACS 8.6 and ACS 8.7. Older versions of ACS software required the StorageWorks Large LUN Utility to be run for a host to access more than eight LUNs per controller port or controller port pair, if using redundant controllers. The Large LUN Utility is located in the Platform Kit that ships with the MA8000 product.

LUN 0 needs to be available and not reserved for Large LUN to function properly. Therefore it is best not to use LUN 0 for a cluster disk. SCSI-3 mode is preferred since LUN 0 is the Command Console LUN (CCL) and is not available.

Note that since mount points are not supported, cluster disks must have drive letters associated with them. Therefore a cluster is able to support a maximum of 22 cluster disks assuming drive letters A, B, C, and D are not used for cluster disks.

For maximum availability of LUNs, use the following configurations:

- Microsoft Windows NT Server 4.0, Enterprise Edition

    — SCSI-3 mode with CCL always enabled

    — LUNs beginning with D1 through D64

- Windows 2000 Advanced Server

    — SCSI-3 mode with CCL always enabled

    — LUNs beginning with D1 through D64

Configuring a Storageset

To configure a storageset:

NOTE: If errors are encountered while creating storagesets, see the installation troubleshooting tips in Chapter 6.

1. Select **Storage** in the **Storage** window menu line.

2. Select **Add Virtual Disk** to begin step 1 of the Add Virtual Disk Wizard.

3. Select one of the available RAID level settings.

4. Click **Next** for step 2 of the Add Virtual Disk Wizard.

5. Select the devices you want to include in the virtual disk by clicking on the disks listed in the **Available Storage** window.

6. Click **Next** for step 3 of the Add Virtual Disk Wizard.

7. Use the displayed value for the virtual disk.

   IMPORTANT: All logical partitions based in a RAID set must be in the same cluster group.

8. Click **Next** for step 4 of the Add Virtual Disk Wizard.

9. Enter the virtual disk name.

10. Verify that the **Save controller configuration to virtual disk** check box is selected.

11. Click **Next** for step 5 of the Add Virtual Disk Wizard. Step 5 will display a summary of your selections.

    NOTE: If you are not satisfied with your selections, return to the applicable Wizard step using the **Back** button. When you are satisfied with your choices, click **Finish.**

12. Click **Finish** to create the virtual disk.

    The **Storage** window displays the virtual disk you created. The hourglass on the disk icon indicates that the storageset is being initialized. The drives you used to create the RAID set are highlighted in the **Devices** window.

13. Repeat steps 1 through 12 to create other virtual disks.

14. Wait for the virtual disk initialization to complete before proceeding. (This process may take more time depending on the size and RAID level of the virtual disk created.)

15. Close the **Storage** window.

Setting Up Connections

To set up HBA connections:

NOTE: Perform the following steps on one HBA at a time for an enhanced configuration.

1. Type the following and press **Enter:**

   ```
   show connections
   ```

   NOTE: Each connection must be deleted and reestablished in order to properly configure the HA/F500. Perform the following step if the connections were previously configured.

2. Type the following and press **Enter:**

   ```
   delete connection !NEWCONXX
   ```

   NOTE: The *XX* in !NEWCON*XX* refers to the connection number.

3. Repeat the above CLI command for each connection.

4. Close the CLI window.

5. Reboot the server you are currently working on.

6. Click **Start** on the Windows taskbar after rebooting.

7. Select **Programs, Command Console, CLI Window.**

8. Click **CLI Window** to display the **Connection Selection** dialog box.

9. Select **Serial** and click **OK** to display the **Connect Serial** dialog box.

10. Select the COM port (the port to which the storage controller is connected).

11. Select a baud rate of **9600.**

12. Click **Connect** to display the CLI window.

13. Press **Enter** to initialize communications.

14. Type the following command and press **Enter:**

    show connections

    NOTE: There should only be one connection, if you have a basic configuration. If not, delete all connections, reboot the server, and repeat the process.

15. Type the following command and press **Enter** to rename the single connection to make it available to the cluster.

    NOTE: The *XX* in !NEWCON*XX* refers to the connection number. Use a naming convention that will be easy to understand and remember. For example: NODE1T where Node1 is the name of the server. T stands for the top storage controller in the storage subsystem. T assumes that the maintenance cable is plugged in the storage controller that resides in the uppermost slot of the controller slots.

    rename !NEWCON*XX* NODE1T

16. Type the following command and press **Enter:**

    show connections

17. Verify that the new connection name is properly set.

18. Boot the second server. Log in to the domain when the login screen displays.

19. Go back to the CLI window, type the following command and press **Enter:**

    show connections

20. If the connection for the second server is not present, reboot the second server and try again.

21. Type the following command and press **Enter** to rename the connection for the second server**.**

    > **NOTE:** The *XX* in !NEWCON*XX* refers to the connection number. Use a naming convention that will be easy to understand and remember. For example: NODE2T where Node2 is the name of the server. T stands for the top storage controller in the storage subsystem. T assumes that the maintenance cable is plugged in the storage controller that resides in the uppermost slot of the controller slots.

    ```
    rename !NEWCONXX NODE2T
    ```

22. Type the following and press **Enter:**

    ```
    show connections
    ```

23. Verify that the new connection name is properly set.

Verifying Access

To verify access to a storageset:

1. Be sure that the second server is powered down.

2. Click **Start** on the Windows taskbar.

3. Select **Programs, Command Console, CLI Window.**

4. Click **CLI Window** to display the **Connection Selection** dialog box.

5. Select **Serial** and click **OK** to display the **Connect Serial** dialog box.

6. Select the COM port (the port to which the storage controller is connected).

7. Select a baud rate of **9600.**

8. Click **Connect** to display the CLI window.

9. Press **Enter** to initialize communications.

10. Type the following and press **Enter:**

    ```
    show units full
    ```

11. Type the following command and press **Enter:**

    `set unit name disable_access_path=ALL`

    NOTE: The ALL access must be removed and set to NONE before adding the connection names defined earlier.

12. Type the following command and press **Enter:**

    `show units full`

13. Verify that the access is now set to **NONE.** If not, repeat the command.

    NOTE: Access will now consist of the connection names defined earlier.

14. Type the following command and press **Enter:**

    `set unit name Enable_Access_Path=connection name`

    Example: `set d0 Enable_Access_Path=NODE1T`

15. Repeat step 14 for the second server connection name:

    Example: `set d0 Enable_Access_Path=NODE2T`

16. Type the following command and press **Enter:**

    `show units full`

17. Verify that the connection names have been properly set. If not, repeat the commands.

18. Repeat steps 11 through 16 for each logical unit to be used in the cluster configuration.

19. Close the CLI window.

20. Reboot the server.

Volume Creation

To create storageset volumes:

1. Run Disk Administrator or Disk Management from one server. Make sure the second server is powered off.

   NOTE: To prevent Windows NT or Windows 2000 from reassigning the local server disks to different drive letters, make them "sticky" by assigning them the drive letters they currently have. This will also prevent drive C: from being renamed to a higher letter.

2. Create a volume on each storageset (virtual disk).

   NOTE: Create only one volume per storageset.

3. Format the volume with the NTFS file system.

4. Assign drive letters. Note the drive letter assignments for later use.

5. Close Disk Administrator or Disk Management.

6. Power off the first server.

7. Power on the second server.

8. Run Disk Administrator or Disk Management from the other server.

9. Assign the same drive letters to the storageset volumes that you assigned on the first server (step 4). The drive letters may not be applied until the server restarts.

   IMPORTANT: The drive letter assignments for the storagesets must be the same on both servers.

10. Restart both servers.

11. After both servers restart, open Disk Administrator or Disk Management on both servers and verify that the drive letters are configured correctly.

Cluster Installation and Verification

The following sections install and verify the cluster installation.

Installing the Cluster Software

Install and configure Microsoft cluster software on the host servers as described in the manual provided with Microsoft Windows NT Server 4.0, Enterprise Edition, Windows 2000 Advanced Server, or Windows 2000 Datacenter Server.

Verifying Creation of the Cluster

To verify the creation of the cluster:

1. Type the following CLI command and press **Enter** to shut down your storage subsystem.

   ```
   shutdown this_controller
   ```

2. Shut down and power off both servers.

3. Power off the storage subsystem and then power it back on.

4. Power both servers back on.

5. From the Windows desktop on either clustered server, select **Start, Programs, Administrative Tools (Common), Cluster Administrator.**

6. Enter the name or IP address of the cluster when you are prompted for **Cluster Name.** If you run Cluster Administrator on a cluster node, enter a period and the cluster is found automatically.

   If the cluster has been created correctly, the computer names of both cluster nodes appear on the left side of the **Cluster Administrator** window.

7. If the cluster is not working correctly, refer to the installation troubleshooting tips in Chapter 6.

Verifying Node Failover

> **IMPORTANT:** Do not run any client activity while testing failover events.

To verify failover of a cluster node:

1. From the desktop on both servers, select **Start, Programs, Administrative Tools (Common), Cluster Administrator.**

2. When you are prompted for the **Cluster Name** or **Server Name,** enter the name or IP address of the cluster.

3. Be sure that all cluster resources and cluster groups are online. Verify that some of the cluster groups are owned by the server you will be powering off so that a failure event will result in a failover of cluster groups.

4. Power off the server mentioned in step 3.

   Within several seconds, Cluster Administrator running on the surviving node should bring online all the predefined resources and groups that were previously owned by the powered-off server. If, after a minute, nothing appears to have occurred, refresh the screen by selecting **Refresh** (or by pressing the **F5** key).

5. If failover is not working correctly, refer to the installation troubleshooting tips in Chapter 6.

6. Power on the powered-off server to continue.

Verifying Network Client Failover

Now that you have verified that the server is correctly running as a cluster node, the next step is to verify that network clients can interact with the cluster.

To verify network client failover:

1. Be sure that both cluster nodes are running and verify through Cluster Administrator that all groups and resources are online.

2. For each hard disk in the shared storage, Microsoft cluster software automatically creates a cluster group that consists of a single resource, the disk drive. Using Cluster Administrator, add an unused IP address as another resource to one of these groups. (Do **not** use the Cluster Group.) Bring the newly created IP resource online.

3. Open a Command Prompt window on a network client machine.

4. Be sure that the network client can access the IP address resource. Regardless of whether you are using WINS or DHCP, you can execute the PING command to check the connection.

   From the network client, execute a PING command using the cluster IP address as the argument. The client has successfully accessed the IP address resource if you get a response similar to:

   `Reply from `*`IP Address`*`: bytes=`*`xx`*` time=`*`xxxms`*` TTL=`*`xx`*`.`

   The client has not successfully accessed the cluster resource if you get a response of:

   `Reply from `*`IP Address`*`: Destination host unreachable`

5. Following the successful completion of the PING command, use Cluster Administrator to perform a manual failover of the cluster group that contains the IP address resource.

6. After the manual failover completes, execute the PING command again.

   As soon as the other node brings the cluster group online, a response similar to the one noted in step 4 should be returned. If the client successfully accessed the failed-over IP address resource, the cluster is working. If the client was unsuccessful, either the cluster group was not configured correctly, the failover did not occur, or the PING command was performed before the failover activity completed.

7. If network client failover is not working correctly, refer to the installation troubleshooting tips in Chapter 6.

8. If you want to verify a more extreme case, rather than fail over the IP address resource, power off the primary cluster node and verify that the resource fails over to the other node.

Setting Up Cluster Groups and Cluster Resources

The clustering hardware is now set up and configured. The Microsoft operating system with Microsoft cluster software is installed on both servers. You have done minimal validation of the cluster and everything is working as planned. Now it is time to set up the applications to use the capabilities of clustering.

If you have defined your cluster needs correctly and determined how to fit the applications and environment into a cluster, then configuring the applications will be a straightforward task.

Although details of these procedures are beyond the scope of this guide, several documents are available from both Microsoft and Compaq to assist you with these final steps.

The best source of information concerning the steps to configure cluster groups and cluster resources is the *Microsoft Cluster Server Administrator Guide,* which provides detailed setup information.

Compaq has worked directly with several application vendors throughout the development of ProLiant Clusters. As a result of these efforts, Compaq has provided a number of integration documents to assist you with installing these applications in a ProLiant Cluster environment.

Integration documents are available at www.compaq.com/highavailability.

Setting Up Dual Storage Controllers

This section describes how to convert a system from a single-adapter, single-storage controller configuration to a single-adapter, dual-storage controller configuration. Dual controller configuration allows failover to occur between the controllers in the event of a controller failure.

Hardware Setup

To set up the hardware:

1. Click the **Start** button on the taskbar.

2. Select **Program, Command Console, CLI Window.**

3. Click **CLI Window** to display the **Connection Selection** dialog box.

4. Select **Serial** and click **OK** to display the **Connect Serial** dialog box.

5. Select the COM port (the port to which the storage controller is connected).

6. Select a baud rate of **9600.**

7. Click **Connect** to display the CLI window.

8. Press **Enter** to initialize communications.

9. Type the following command and press **Enter:**

   ```
   shutdown this_controller
   ```

10. Close the CLI window.

11. Shut down and power off both host servers.

12. Power off the storage subsystem.

13. Install the second storage controller into the empty slot (see Figure 3-4).
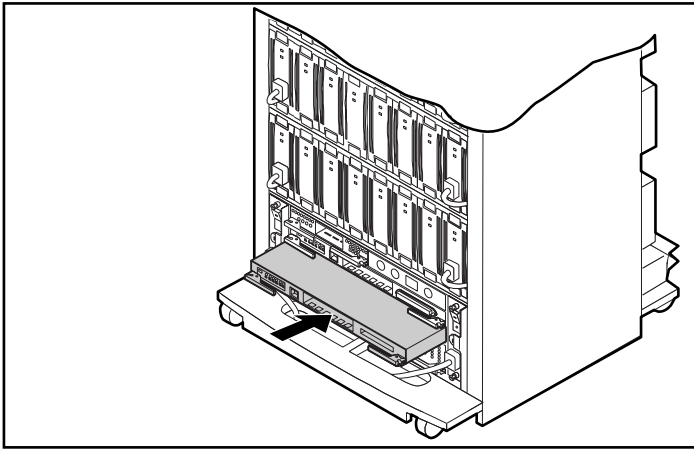
Figure 3-4: Installing the second storage controller (RA8000 shown)

14. Install the PCMCIA card into the controller (see Figure 3-5).

15. Connect port 1 (the left port) of the second controller to the Fibre Channel switch.
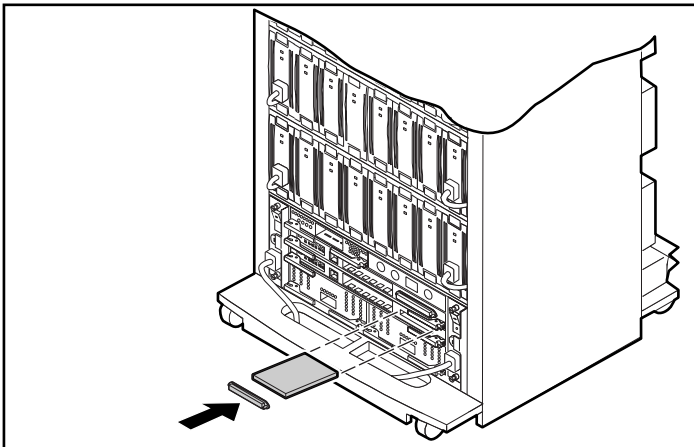
Figure 3-5: Installing the PCMCIA into the second storage controller (RA8000 shown)

16. Power on the storage subsystem.

17. Power on both host servers.

Configuring Controllers for Failover Mode

NOTE: To use two storage controllers and one HBA in each server, you need to configure the controllers to Transparent Failover mode.

1. Click the **Start** button on the taskbar.

2. Select **Program, Command Console, CLI Window.**

3. Click **CLI Window** to display the **Connection Selection** dialog box.

4. Select **Serial** and click **OK** to display the **Connect Serial** dialog box.

5. Select the COM port (the port to which the storage controller is connected).

6. Select a baud rate of **9600.**

7. Click **Connect** to display the CLI window.

8. Press **Enter** to initialize communications.

9. Type the following command and press **Enter:**

   ```
   set failover copy=this_controller
   ```

10. Wait two minutes for the controllers to reset.

11. Verify that both controllers are in dual redundant configuration (failover) mode by entering the following commands and viewing the results:

    ```
    show this_controller
    ```

    ```
    show other_controller
    ```

12. Close the CLI window.

Verifying the System

Restart both servers and inspect the cluster and the event logs of the hosts to verify that the system is functioning without errors.

Setting Up a Single Cluster with an Additional Storage Subsystem (Single Controller Configuration)

1. Set up the cluster with the first storage subsystem according to the instructions in this chapter starting on page 3-5.

2. Follow the cluster installation instructions in this chapter to set up the second, third, or fourth storage subsystem, with the **important exceptions** included in Table 3-6.

Table 3-6: Setting Up a Single Cluster with an Additional Storage Subsystem (Single Controller Configuration)

Chapter 3 Section Heading	Page	Action
Setting Up the Servers	3-5	Omit
Updating Switch Firmware	3-10	Omit
Combining Switch Models in a Cluster Configuration	3-11	Omit
Installing the StorageWorks Command Console Client	3-11	Omit
Installing the Adapter Driver	3-14	Omit
Running the Fibre Channel Fabric Setup Utility	3-16	Omit
Cluster Installation and Verification	3-26	Omit
Setting Up Dual Storage Controllers	3-29	Omit

NOTE: Refer to Chapter 5 for information on adding shared storage.

Setting Up a Single Cluster with an Additional Storage Subsystem (Dual Controller Configuration, Transparent Failover)

1. Set up the cluster with the first storage subsystem according to the instructions in this chapter starting on page 3-5.

2. Follow the cluster installation instructions in this chapter to set up the second, third, or fourth storage subsystem, with the **important exceptions** included in Table 3-7.

Table 3-7: Setting Up a Single Cluster with an Additional Storage Subsystem (Dual Controllers, Transparent Failover)

Chapter 3 Section Heading	Page	Action
Setting Up the Servers	3-5	Omit
Updating Swtich Firmware	3-10	Omit
Combining Switch Models in a Cluster Configuration	3-11	Omit
Installing the StorageWorks Command Console Client	3-11	Omit
Installing the Adapter Driver	3-14	Omit
Running the Fibre Channel Fabric Setup Utility	3-16	Omit
Cluster Installation and Verification	3-26	Omit

NOTE: Refer to Chapter 5 for information on adding shared storage.

Multiple Cluster Configurations for Switches

This section contains instructions for configuring multiple clusters using the same storage subsystem. This section will refer to steps that can be found in Chapters 3 and 4 of this guide.

Key Considerations

When setting up multiple cluster configurations, be aware of the following key considerations:

- Maximum configuration in a switch environment
- Virtual disk naming conventions
- Server-to-storage connection paths

Maximum Configuration in a Switch Environment

A switch environment is currently capable of supporting four clusters (two servers per cluster) and four single storage controllers in four separate storage subsystems (or four storage controller pairs in four separate storage subsystems). This also enables a single cluster to have up to four separate storage subsystems or four clusters to share a single storage subsystem. Considerations are necessary when setting up the first cluster if the first storage subsystem will be shared by more than one cluster in the future. Considerations include leaving disks available for additional clusters or installing expansion cabinets for additional storage.

Virtual Disk Naming Conventions

Table 3-8 shows the naming conventions used with the Large LUN Utility.

Table 3-8: Storage Controller Configurations

	Unit Offset	Drive Labels
Cluster 1	0	D1-D64
Cluster 2	100	D101-D164

Server-to-Storage Connection Paths

Server-to-storage connection paths must be defined when using a switch environment because additional clusters will be able to see all drives on a switch. The connection path setting helps segregate the logical units among the clusters.

Single Storage Controller Configurations

The following sections provide information on setting up single storage controller configurations.

Setting Up an Additional Cluster with a Single Storage Subsystem

To set up an additional cluster with a single storage subsystem:

1. Set up the first cluster according to the instructions starting on page 3-5.

2. Follow the cluster installation instructions in this chapter to set up the second, third, or fourth cluster with the **important exceptions** included in Table 3-9.

Table 3-9: Setting Up the Second Cluster
(Single Controller Configuration, Single Storage Subsystem)

Chapter 3 Section Heading	Page	Action
Set Controller to No Failover	3-12	Omit
Disabling the Command Console LUN	3-13	Omit
Setting Port Topology	3-14	Omit
Discovering Devices	3-14	Execute section if additional drives are added.

continued

**Table 3-9: Setting Up the Second Cluster
(Single Controller Configuration, Single Storage Subsystem)** *continued*

Chapter 3 Section Heading	Page	Action
Verifying Controller Properties	3-17	Omit
Verifying Access	3-21	Step 12: Delete only the connections that have !NEWCONXX. Leave the connection paths you renamed previously.
		Only add the connection names that need access to the new logical units.
Setting Up Dual Storage Controllers	3-29	Omit

Setting Up Additional Clusters and Storage Subsystems

NOTE: When installing additional clusters and additional storage subsystems, minimize rebooting the existing cluster nodes because connection paths will be established on the new storage subsystem for all servers that reboot and rescan for all available logical units. This can become cumbersome when trying to establish permanent connection path names for the new servers.

1. Set up the first cluster according to the instructions starting on page 3-5.

2. Follow the cluster installation instructions in this chapter to set up the second, third, or fourth cluster with the **important exceptions** included in Table 3-10.

Table 3-10: Setting Up Additional Clusters and Storage Subsystems

Chapter 3 Section Headings	Page	Action
Setting Up Dual Storage Controllers	3-29	Omit

Setting Up Dual Redundant Controllers (Transparent Failover Configurations)

The following text provides instructions for setting up a multiple cluster configuration with dual redundant controllers (transparent failover).

Setting Up an Additional Cluster with a Single Storage Subsystem

To set up an additional cluster with a single storage subsystem:

1. Set up the first cluster according to the instructions starting on page 3-5.

2. Follow the cluster installation instructions in this chapter to set up the second, third, or fourth cluster with the **important exceptions** included in Table 3-11.

Table 3-11: Setting Up the Second Cluster (Dual Controller Configuration)

Chapter 3 Section Heading	Page	Action
Setting Controller to No Failover	3-12	Omit
Disabling the Command Console LUN	3-13	Omit
Setting Port Topology	3-14	Omit
Discovering Devices	3-14	Execute section if additional drives are added.
Verifying Controller Properties	3-17	Omit
Verifying Access	3-21	Step 12: Delete only the connections that have !NEWCON*XX*. Leave the connection paths you renamed previously.
		Only add the connection names that need access to the new logical units.

Setting Up Additional Clusters and Storage Subsystems

> **NOTE:** When installing an additional clusters and additional storage subsystems, minimize rebooting the existing cluster nodes because connection paths will be established on the new storage subsystem for all servers that reboot and rescan for all available logical units. This can become cumbersome when trying to establish permanent connection path names for the new servers.

Set up the additional clusters according to the instructions starting on page 3-5.

Setting Up Dual Redundant Controllers (Multibus Failover)

The following text provides instructions for setting up a multiple cluster configuration with dual redundant controllers (multibus failover).

Setting Up an Additional Cluster with a Single Storage Subsystem

To set up an additional cluster with a single storage subsystem:

1. Set up the first cluster according to the instructions in Chapters 3 and 4 of this guide.

2. Follow the cluster installation instructions in Chapters 3 and 4 to set up the second, third, or fourth cluster with the **important exceptions** included in Table 3-12.

Table 3-12: Setting Up the Second Cluster (Single Controller Configuration)

Chapter 3 and 4 Section Heading	Page	Action
Setting Controller to No Failover	3-12	Omit
Disabling the Command Console LUN	3-13	Omit
Setting Port Topology	3-14	Omit
Discovering Devices	3-14	Execute section if additional drives are added.
Verifying Controller Properties	3-17	Omit
Verifying Access	3-21	Step 12: Delete only the connections that have !NEWCONXX. Leave the connection paths you renamed previously.

Only add the connection names that need access to the new logical units. |
Setting Storage Controllers to Multibus Failover	4-5	Omit
Establishing a Redundant Path for the Microsoft Cluster	4-11	Omit
For Switches Only: steps 1–8		

Setting Up Additional Clusters and Storage Subsystems

NOTE: When installing additional clusters and additional storage subsystems, minimize rebooting the existing cluster nodes because connection paths will be established on the new storage subsystem for all servers that reboot and rescan for all available logical units. This can become cumbersome when trying to establish permanent connection path names for the new servers.

Set up the additional clusters according to the instructions in Chapters 3 and 4 of this guide.

4

Secure Path Installation Procedures

Preinstallation Instructions

Before installing Secure Path, verify that the hardware and software are appropriate to this installation. The following section provides more details about the hardware and software.

Checking the Hardware and Software

Please verify that you have received the following hardware (not installed in Chapter 2) and the following software:

NOTE: Depending on the configuration, you may have either Fibre Channel hubs or Fibre Channel switches.

Table 4-1: Storage Subsystem with Microsoft Clusters

Quantity	Description
1	Secure Path for Windows
2	Host bus adapter
1	12-port Fibre Channel hub OR 7-port Fibre Channel hub

continued

Table 4-1: Storage Subsystem with Microsoft Clusters *continued*

Quantity	Description
1	8-port Fibre Channel switch OR 16-port Fibre Channel switch
1	Fibre Channel interconnect hardware

If you are missing any component required for your Secure Path environment, contact your local sales representative or call 1-800-OK-COMPAQ.

Examining the Current Configuration

If you have just performed the steps outlined in Chapter 2 or Chapter 3, then the basic cluster configuration conforms to the requirements listed in this section. You may skip this and the next section and proceed directly to "Installing Secure Path Software." Otherwise, be sure that the existing single path configuration conforms to Secure Path requirements. The requirements include:

1. Verifying that there is a serial connection to the storage subsystem and that you can communicate to it through SWCC

2. Launching the event log viewer and verifying that the Fibre Channel Setup Utility is installed and that it reports the expected number of logical units

3. Checking the event log viewer to ensure that there are no error events reported by the host bus adapter or Fibre Channel Setup Utility

4. Verifying that the operating system (boot) disk is not part of the storage subsystem

5. Verifying that none of the volume sets uses software RAID or extended volumes

6. Verifying that each server has the TCP/IP protocol installed and that each server is available on the network by pinging it

Preparing an Existing Storage Subsystem for Secure Path Operation

Perform the following steps before proceeding to installation procedures if you have an existing storage subsystem that is currently being used in a production environment and plan to reconfigure for Secure Path operation.

1. Follow normal procedures to back up the data stored on all drives configured on the storage subsystem.

> △ **CAUTION:** If you currently have a storage subsystem in a production environment that is being converted to Secure Path operation, be sure that all end users have logged off the server and that all I/O to the storage subsystem has ceased before proceeding.

2. Before installing Secure Path, refer to Table 4-2 for the maximum logical unit limits that Secure Path supports.

Table 4-2: Secure Path LUN Maximums

Maximum LUNs per Cluster without Large LUN	Maximum LUNs per Cluster with Large LUN	Number of Clusters	Topology
16	64	1	Loop
8	64	2	Loop
16	64	Up to 4	Switch (non-DT)

> △ **CAUTION:** Secure Path only recognizes the maximum logical units listed in Table 4-2 after the software is installed.

NOTE: Offsets may be used, but to properly restrict access, use the selective storage presentation. Units should begin with D1 and continue consecutively.

Installing Secure Path Software

The following sections explain how to install the Secure Path software.

Overview

The following sections describe the software configuration procedures required for the Secure Path storage environment:

- Installing the Secure Path software on the host servers

- Establishing a serial link to the storage subsystem

- Setting storage controllers to multibus failover mode

- Preferring the paths of the storagesets (units) to the storage controllers

After performing these procedures in sequence, the software configuration of your Secure Path storage environment will be complete.

Description of Secure Path Software

Secure Path for Windows consists of a kernel mode driver responsible for directing I/O to the desired path and for changing paths whenever the driver detects failure in a redundant path.

Secure Path for Windows is managed by a client/server management application that requires TCP/IP to be installed in the Windows NT or Windows 2000 server attached to the storage subsystem where the Secure Path agent is installed, and on the management station on which the Secure Path Manager graphical user interface is installed.

The Secure Path user interface and agent (client/server) may be installed in the same server, as long as the agent is installed on the server that is attached to the storage subsystem to be managed.

Configuring the Storage Subsystem for Secure Path Operation

This section describes how to configure the storage subsystem controllers for a Secure Path environment, which includes:

- Setting the controllers to the multibus failover mode

- Preferring (specifying) which storage controller the I/O of each disk will be assigned to upon system boot

Setting Storage Controllers to Multibus Failover

Secure Path operation requires that the storage controllers be configured for a multiple bus failover mode. To configure a storage controller for multiple bus failover mode:

1. Click **Start** on the taskbar from the maintenance server.

2. Select **Programs, Command Console, CLI Window.**

3. Click **CLI Window** to display the **Connection Selection** dialog box.

4. Select **Serial** and click **OK** to display the **Connect Serial** dialog box.

5. Select the COM port (the port to which the storage controller is connected).

6. Select a baud rate of **9600.**

7. Click **Connect** to display the CLI window.

   NOTE: You may need to press **Enter** to initialize communications before entering CLI commands.

8. Type the following command and press **Enter:**

   ```
   set nofailover
   ```

   The "other" controller will shut down and must be manually restarted by momentarily depressing the reset button on the controller front panel. Wait two minutes for the controller to boot before proceeding.

9. Type the following and press **Enter:**

   ```
   set multibus copy=this_controller
   ```

   The controllers will restart in multibus mode. Wait two minutes for the controllers to restart.

10. Verify that both controllers are configured for multibus failover mode by issuing the following commands and viewing the results:

   ```
   show this_controller
   ```

   ```
   show other_controller
   ```

The controllers are now configured for multibus operation.

Preferring Storage Unit Paths

To complete the multibus configuration setup, you must assign storage units to one of the controllers to specify which controller is used to access units at system boot time. The preferred_path unit attribute assigns units to either "this" or the "other" controller. In effect, this procedure specifies on which path (controller and/or host bus adapter) the I/O will travel.

It is recommend that you initially balance the available storagesets between the buses. As storage demands are defined and individual drive throughput requirements are understood, adjustments to the disk I/O path configuration may be made using the SANworks Secure Path Manager.

1. From the CLI window, enter the following command to obtain a list of all units defined in the storage subsystem:

   ```
   show units
   ```

2. Enter the following command to specify preferred_path for units:

   ```
   set unit# preferred=this_controller
   ```

   or

   ```
   set unit# preferred=other_controller
   ```

3. Verify that the preceding changes were made by entering the following command:

```
show units full
```

You have completed the software configuration required to support the Secure Path environment. Proceed to the next section to cable the second path. Then you will be ready to monitor and manage Secure Path activity using the SANworks Secure Path Manager.

Installing Hardware

This section provides the procedures for installing and terminating a second I/O path between a storage subsystem and the clustered servers, where a single I/O path currently exists.

NOTE: Follow normal procedures to power off the server prior to cabling.

Hardware Configuration Overview

Configuring the hardware components consists of three main tasks to be performed in sequence, as described in the following sections:

1. Installing the second host bus adapter
2. Cabling the hardware components
3. Verifying the hardware configuration

Establishing a Redundant Path for the Microsoft Cluster

The following sections provide information on establishing a redundant path for hubs and switches.

For Hubs Only

To establish two individual Fibre Channel loops between clustered host servers and a storage subsystem where one loop currently exists:

IMPORTANT: Before proceeding, be sure that the Secure Path software is installed.

1. Open the CLI window and press **Enter** to initialize communication.

2. Enter the following commands:

   ```
   shutdown other_controller
   shutdown this_controller
   ```

3. Shut down and power off the servers and external storage subsystem.

4. Install the second host bus adapter in each server (refer to Chapter 2).

5. Connect the second hub to the second host bus adapter in each server and to the second storage controller.

   IMPORTANT: You must use only one set of ports in the controller pair. It is recommended that you use one port (the left port) on each controller.

6. Power up the storage subsystem.

7. Power up the server with the maintenance terminal.

8. Log in to the domain.

9. Select **Start, Programs, Command Console, CLI Window.**

10. Click **CLI Window** to display the **Connection Selection** dialog box.

11. Select **Serial** and click **OK** to display the **Connect Serial** dialog box.

12. Select the COMM port (the port to which the storage controller is connected).

13. Select a baud rate of **9600.**

14. Click **Connect** to display the CLI window.

15. Press **Enter** to initialize communications.

16. Type the following command and press **Enter:**

    ```
    show connections
    ```

17. There should only be one new connection (!NEWCON*XX*) for Server 1.

     NOTE: The "*XX*" in "!NEWCON*XX*" refers to the connection number.

18. Type the following command and press **Enter** to rename the single connection to be available for the redundant path:

     NOTE: The "*XX*" in "!NEWCON*XX*" refers to the connection number.

     NOTE: Use a naming convention that will be easy to understand and remember. For example: NODE1B where Node 1 is the server and B stands for the bottom storage controller in the storage subsystem. B assumes that the maintenance cable is plugged into the other controller.

     ```
     rename !NEWCONXX NODE1B
     ```

19. Type the following command and press **Enter:**

     ```
     show connections
     ```

20. Verify that the new connection name is properly set.

21. Boot the second server. Log in to the domain when the log-in screen appears.

22. Go back to the CLI window, type the following command, and press **Enter:**

     ```
     show connections
     ```

23. The second server connection should now be present as well. If not, reboot the second server and try again.

24. Type the following command and press **Enter** to rename the second server connection:

   NOTE: The "*XX*" in "!NEWCON*XX*" refers to the connection number.

   NOTE: Use a naming convention that will be easy to understand and remember. For example: NODE2B where Node 2 is the server. B stands for the top storage controller in the storage subsystem. B assumes that the maintenance cable is plugged into the other controller.

   When you have two HBAs, you may want to use a naming convention that is more detailed to the equipment. For example,

   rename !NEWCON*XX* N1A1P1

   N1 is Node 1, A1 is adapter 1, and P1 is port 1 on the controller.

   ```
   rename !NEWCONXX NODE2B
   ```

25. Type the following and press **Enter:**

   ```
   show connections
   ```

26. Verify that the new connection name is properly set. The new connection names also need to be added.

27. Type the following and press **Enter:**

   ```
   set unit name enable_access_path=connection name
   ```

   Example: `set d0 enable_access_path=NODE1B`

28. Repeat the preceding command for the second server connection name:

   Example: `set d0 Enable_Access_Path=NODE2B`

29. Type the following command and press **Enter:**

   ```
   show units full
   ```

30. Verify that the connection names are properly set. If not, repeat the commands.

31. Repeat steps 27 through 30 for each logical unit to be used in the cluster configuration.

32. Close the CLI window.

Secure Path is now properly prepared and cabled.

For Switches Only

To establish two individual Fibre Channel fabrics between clustered host servers and a storage subsystem where only one fabric currently exists:

IMPORTANT: Before proceeding, be sure that the Secure Path software is installed.

1. Open the CLI window and press **Enter** to initialize communication.

2. Enter the following commands:

   ```
   shutdown other_controller
   shutdown this_controller
   ```

3. Shut down and power off the servers and external storage subsystem.

4. Install the second host bus adapter in each server (refer to Chapter 3).

5. Set up the second switch. Refer to "Setting Up the Switch" in Chapter 3. Complete these steps before continuing to step 6.

6. Connect the second switch to the second host bus adapter in each server and to the second storage controller.

7. Power up the storage subsystem.

8. Power up the server with the maintenance terminal.

9. Log in to the domain.

10. Select **Start, Programs, Command Console, CLI Window.**

11. Click **CLI Window** to display the **Connection Selection** dialog box.

12. Select **Serial** and click **OK** to display the **Connect Serial** dialog box.

13. Select the COMM port (the port to which the storage controller is connected).

14. Select a baud rate of **9600.**

15. Click **Connect** to display the CLI window.

16. Press **Enter** to initialize communications.

17. Type the following command and press **Enter:**

    ```
    show connections
    ```

18. Verify that there is only one new connection (!NEWCON*XX*) for Server 1.

    NOTE: The "*XX*" in "!NEWCON*XX*" refers to the connection number.

19. Type the following command and press **Enter** to rename the single connection to be available for the redundant path:

    NOTE: The "*XX*" in "!NEWCON*XX*" refers to the connection number.

    NOTE: Use a naming convention that will be easy to understand and remember. For example: NODE1B where Node 1 is the server and B stands for the bottom storage controller in the storage subsystem. B assumes that the maintenance cable is plugged into the other controller.

    ```
    rename !NEWCONXX NODE1B
    ```

20. Type the following command and press **Enter:**

    ```
    show connections
    ```

21. Verify that the new connection name is properly set.

22. Boot the second server. Log in to the domain when the log-in screen appears.

23. Go back to the CLI window, type the following command, and press **Enter:**

    ```
    show connections
    ```

24. The second server connection should now be present as well. If not, reboot the second server and try again.

25. Type the following command and press **Enter** to rename the second server connection:

    NOTE: The "*XX*" in "!NEWCON*XX*" refers to the connection number.

    NOTE: Use a naming convention that will be easy to understand and remember. For example: NODE2B where Node 2 is the server. B stands for the top storage controller in the storage subsystem. B assumes that the maintenance cable is plugged into the other controller.

    When you have two HBAs, you may want to use a naming convention that is more detailed to the equipment. For example,

    rename !NEWCON*XX* N1A1P1

    N1 is Node 1, A1 is adapter 1, and P1 is port 1 on the controller.

    ```
    rename !NEWCONXX NODE2B
    ```

26. Type the following and press **Enter:**

    ```
    show connections
    ```

27. Verify that the new connection name is properly set. The new connection names also need to be added.

28. Type the following and press **Enter:**

    set *unit name* enable_access_path=*connection name*

    Example: set d0 enable_access_path=NODE1B

29. Repeat the preceding command for the second server connection name:

    Example: set d0 Enable_Access_Path=NODE2B

30. Type the following command and press **Enter:**

    ```
    show units full
    ```

31. Verify that the connection names are properly set. If not, repeat the commands.

32. Repeat steps 28 through 31 for each logical unit to be used in the cluster configuration.

33. Close the CLI window.

Secure Path is now properly prepared and cabled.

Verifying the Secure Path Hardware Configuration

After system bootup, check the operating system event log for successful start events for the RaiDisk and HszDisk.

Using Secure Path Manager

This section describes how to use Secure Path Manager to monitor and manage a Secure Path for a Windows NT or Windows 2000 environment.

NOTE: This section assumes that storage subsystem storagesets have already been configured, using SWCC, and that the drives have been partitioned and formatted with Windows Disk Administrator. These procedures are described in the getting started guide, which shipped with your subsystem.

Secure Path Manager is a Graphical User Interface (GUI) utility that:

- Reports the status of the two paths

- Facilitates balancing I/O between the two bus paths

- Reports disk status (path assignment, failover, and failback activity)

- Enables (manual) drive failback upon path restoration

Compaq recommends that the Secure Path Management application remain active (or minimized) to provide continuous Secure Path status monitoring. To monitor and manage a Secure Path environment using Secure Path Manager as described in the following sections:

1. From the **Start** menu, select **Programs, SANworks.**

2. Select the SPM application icon.

3. Log in by entering the server name and password previously assigned in the "Installing the Secure Path Software" section.

   IMPORTANT: If you make changes to the Secure Path Agent Configuration, the changes will not go into effect until you stop and restart the Secure Path Agent Service from the control panel.

Path and Drive Status Monitor

The two paths (Path 0 and Path 1) are displayed in green by the SecurePath Manager when both paths (host bus adapters, cabling, and controllers) are functioning normally. The Manager keeps track of the primary path (bus assignment) for each disk in the storage subsystem. When operating normally, the Manager displays each disk on the primary path to which it has been assigned as a disk icon, shaded yellow and gray. If a drive letter has been assigned to the disk, it will be displayed above the disk icon.

Determining Disk Identity

Drives configured in a multipath environment may be identified in three ways. If you move the cursor over a disk icon, Secure Path Manager will display the various identities of a drive/port/bus/target/LUN and disk number. You may also view this information by right-clicking the mouse on the disk icon to launch the properties dialog for the drive.

The port/bus/target/LUN information refers to the physical identity of the drive's corresponding storageset (unit) as designated by the storage subsystem. The disk number refers to the number assigned to a drive by the Windows Disk Administrator, and the drive letter appearing above the drive icon is assigned to that partition (if one exists) by the system administrator. If a drive has more than one partition, the disk number also appears above the drive. This drive information should allow you to quickly map the storageset to operating system identity and determine which storagesets are currently serviced through each path.

Assigning New Primary Paths to Drives

To assign a new primary path to a drive:

1. Use the left mouse key to select the icon of the disk that is to receive a new primary path assignment (the icon becomes a null-circle when selected).

2. Drag the disk icon from its current primary path to the alternate path displayed on the screen (the icon turns gray while in transition, and the cursor changes to a squared arrow).

3. Drop the disk icon anywhere you see the squared arrow along the "new" primary (formerly alternate).

   When the primary path reassignment of a disk completes, its icon is displayed in original form, shaded gray and yellow, on the new path. (It may appear above or beneath the path line, depending on your exact placement of the mouse). Repeat this procedure for each disk that is to be assigned a new primary path.

The Manager does not permit a new primary path assignment to a failed path. If an attempt is made to move a drive to a failed path, the Manager returns the drive to the original path. A pop-up window also appears stating that the failback was not successful.

NOTE: The order and spacing of the disk icons displayed on the paths is refreshed every 90 seconds, and can be refreshed any time by using the **View/Refresh** pull-down menu of the Manager window or by pressing the **F5** key.

Balancing the I/O Load Between Paths

As the storage demands of the Secure Path environment are defined and individual drive throughput requirements are understood, Compaq recommends that you evenly balance the disks generating the highest I/O loads between the two paths to maximize overall throughput. The Manager may be used to statically load-balance the Secure Path configuration by following the procedure below:

1. Identify "hot" drives which are the drives that consistently experience the greatest I/O load while running workloads typical of the production environment. Enable operating system disk performance statistics, if you have not done so already, by issuing *diskperf -y* from a command window and restarting the system.

   Use Windows NT Performance Monitor to characterize individual drive loading in terms of throughput (I/Os per second) and/or bandwidth (bytes per second), whichever is more appropriate for the application.

2. Note the path assignments of hot drives.

3. Balance the overall I/O load, as much as possible, by evenly distributing (reassigning primary drive path) the hot drives between the two paths.

Run the workload, monitor the operation, and readjust as necessary.

Defining a Persistent Secure Path Storage Subsystem Drive Configuration

When the primary path for a drive is changed using the Manager, the preferred path assignment for the corresponding storage unit on the storage subsystem does not change. If the preferred path is not changed to the new path, the unit reverts to its original preferred path if both the storage subsystem and host server are power cycled together. To make the primary path assignment persistent for those drives you have reassigned with the Manager, reset the preferred path attribute for the corresponding storage unit on the storage subsystem.

1. Use the CLI command `show units` to show the preferred path settings for all units. This command also indicates which controller each storageset is currently online with ("this" or "other").

2. Next, use the CLI command `set unit# preferred=this_controller` or `set unit# preferred=other_controller` to change the preferred path attribute to the appropriate path. For instance, if a storage unit is reported as being "online to other controller" but is preferred to "this" controller, you should change the preferred path attribute to the "other" controller.

3. Repeat this procedure for each storage unit that is online to a path that is not its preferred path. It is not necessary to restart the server or storage subsystem to perform this procedure.

Automatic Failover

When a path fails (Secure Path software detects the loss of drive I/O due to adapter, cable, or controller malfunction), the Secure Path software will:

*   Perform an automatic failover and move the affected drives to the alternate path.

*   Log failover events in the operating system Event Log.

*   Report the path failure through a Windows pop-up message.

*   Reflect the reassignment of the drives to the failover path on the display.

    NOTE: Check the operating system Event Log or entries generated by the Secure Path software to help determine which component of the path has malfunctioned. Look for entries by the HszDisk and RaiDisk drivers.

Automatic Failover Detection and Status Reporting

The Secure Path software continuously monitors the operational status of drives configured on each path. If the Secure Path software detects the failure of an I/O to complete a drive, it will immediately move that drive to its alternate path and reroute outstanding I/O accordingly. Following the occurrence of any drive failure, the Manager reflects the updated Secure Path configuration within its second refresh interval or sooner if the end user presses the **F5** key.

When the Manager discovers the failover of at least one drive, it generates a Windows pop-up message and designates the path as failed. Because the Secure Path software detects path failure through failed I/O operations, only those drives with I/O active at the time of the failure will fail over. Those without active I/O will remain on the failed path by a drag-and-drop operation.

NOTE: For a quick reference of the disk colors and their meaning, select **Legend** from the **View** pull-down menu. The three possible drive states are identified and displayed in color. Refer to the Help files for possible path colors.

Manual Failback and Status Reporting

After a failed path is restored, the disks that had been failed over to an alternate path must be failed back manually, one at a time. As a safeguard, the Secure Path Manager does not automatically fail back drives. Instead, it enables disks to be manually failed back after the administrator validates the integrity of the path.

For a failed path to return to the normal (green) state, the path must be restored and one or both of the following events must occur:

- All of the failed-over disks are failed back to it.

- The server is restarted.

Failback Methods

Failed-over drives may be restored to their primary path using one of the following failback methods:

- Double-click on the disk icon to be failed back.

- From the **Failback** drop-down menu, select the **Failback** option, select the disk, then click **OK.**

- In the Manager Toolbar, click **Failback**, select the disk, then click **OK.**

- Drag and drop each failed-over disk icon to its primary path as follows:

  — Use the left mouse key to select the icon of the disk that is to be failed back to its assigned primary path (the icon will become a null-circle when selected).

  — Drag the disk icon from its current, alternate path to the primary path displayed on the screen (the icon turns gray while in transition, and the cursor changes to a squared arrow).

  — Drop the disk icon anywhere along the disk primary path.

If more than one drive requires failback, repeat these steps until all drives are failed back and the restored path turns green. As each individual drive returns to its primary path, the drive icons will return to their normal yellow and gray color. (The disk icon may appear above or beneath the path line, depending on your exact placement of the mouse.) The color of the failback path will not return to green until all failed-over drives have been restored to their primary path.

NOTE: The order and spacing of the disk icons displayed on the paths can be refreshed using the **View/Refresh** drop-down menu of the **Secure Path Manager** window or by pressing the **F5** key.

Adding New Storagesets with Secure Path

To add a new storageset to a Secure Path configuration:

1. Use the **HSG80 Storage** window to create new storagesets on the storage subsystem.

2. Follow the procedures in the "Configuring a Storageset" section in Chapter 2 to configure a new storageset unit.

3. Use the CLI window to assign a preferred path to the new unit.

4. Use appropriate procedures to add a new volume to the server or cluster.

5. Restart both host servers so that Windows NT or Windows 2000 and the Secure Path software can configure the new unit.

Refer to the *Microsoft Cluster Server Administrator Guide* to include new volumes in the cluster.

Removing a Storageset with Secure Path

To remove a storageset:

1. Use Windows Disk Administrator to delete the partition from the drive to be removed, using appropriate procedures.

2. Shut down the operating system on both host servers.

3. Use SWCC to delete the storagesets on the storage subsystem.

4. Restart both host servers to allow the Secure Path software to configure storage devices.

Setting Up a Single Cluster with an Additional Storage Subsystem (Dual Controller Configuration, Multibus Failover) with Hubs

To set up a single cluster with an additional storage subsystem:

1. Set up the cluster with the first storage subsystem according to the instructions in Chapters 2 and 4 of this guide.

2. Follow the cluster installation instructions in Chapters 2 and 4 of this guide to set up the second, third, or fourth storage subsystem, with the **important exceptions** included in Table 4-3.

Table 4-3: Setting Up a Single Cluster with an Additional Storage Subsystem (Dual Controller Configuration, Multibus Failover) with Hubs

Chapters 2 and 4 Section Heading	Page	Action
Setting Up the Servers	2-3	Omit
Installing the Host Bus Adapter	2-5	Omit
Installing the StorageWorks Command Console Client	2-8	Omit
Installing the Adapter Driver	2-11	Omit
Cluster Installation and Verification	2-22	Omit
Installing Secure Path Software	4-4	Omit
Hardware Configuration Overview	4-7	Omit
Establishing a Redundant Path for the Microsoft Cluster	4-8	Step 2: Perform on all storage subsystems
For Hubs Only		Steps 4: Omit
		Step 5: Connect the top controller to the first hub. Connect the bottom controller to the second hub.

NOTE: Refer to Chapter 5 for information on adding shared storage.

Setting Up a Single Cluster with an Additional Storage Subsystem (Dual Controller Configuration, Multibus Failover) with Switches

To set up a single cluster with an additional storage subsystem:

1. Set up the cluster with the first storage subsystem according to the instructions in Chapters 3 and 4 of this guide.

2. Follow the cluster installation instructions in Chapters 3 and 4 of this guide to set up the second, third, or fourth storage subsystem, with the **important exceptions** included in Table 4-4.

Table 4-4: Setting Up a Single Cluster with an Additional Storage Subsystem (Dual Controller Configuration, Multibus Failover) with Switches

Chapters 3 and 4 Section Heading	Page	Action
Setting Up the Servers	3-5	Omit
Installing the Switch	3-10	Omit steps 4 and 5
Installing the StorageWorks Command Console Client	3-11	Omit
Installing the Adapter Drive	3-15	Omit
Running the Fibre Channel Fabric Setup Utility	3-16	Omit
Cluster Installation and Verification	3-26	Omit
Hardware Configuration Overview	4-7	Omit

continued

Table 4-4: Setting Up a Single Cluster with an Additional Storage Subsystem (Dual Controller Configuration, Multibus Failover) with Switches *continued*

Chapters 3 and 4 Section Heading	Page	Action
Establishing a Redundant Path for the Microsoft Cluster	4-11	Step 2: Perform on all storage subsystems
For Switches Only		Steps 4-5: Omit
		Step 6: Connect the top controller to the first switch. Connect the bottom controller to the second switch.
		Step 7: The maintenance terminal has to be moved to each additional storage subsystem to perform the following steps in the section.
Verifying the Secure Path Hardware Configuration	4-14	Omit

NOTE: Refer to Chapter 5 for information on adding shared storage.

5

Managing the ProLiant Cluster HA/F500

Throughout the life of the cluster, you will encounter the need to improve performance, upgrade hardware components, upgrade software, increase storage capacity, restructure cluster groups, back up the cluster data, or monitor ongoing activities. This chapter describes these management concepts as they pertain to clusters. The chapter also details the utilities and programs used in the ongoing management of the ProLiant Cluster HA/F500.

Cluster Management Concepts

The following sections provide information on various cluster management concepts.

Managing a Cluster Without Interrupting Cluster Services

At some time during the life of the cluster, you may need to perform an operation on a cluster node that requires it to be powered off. Always use Cluster Administrator to fail over (or at a minimum, bring offline) clustered applications before powering off the server.

Managing a Cluster in a Degraded Condition

Due to the high-availability nature of clustering, applications and network clients remain operational even while some cluster components do not. When the cluster is in a degraded condition, it is important to follow this process:

1. Discover what caused the degradation.

2. Determine whether the condition will continue to worsen.

3. Determine how critical it is to repair the problem.

    a. If the problem is considered noncritical, wait until a nonpeak time to service the problem.

    b. If the problem is considered critical, fail over all clustered applications and resources to the other server before servicing the problem.

Managing Network Clients Connected to a Cluster

An important aspect of managing network clients is to inform end users that their applications are now running on a cluster. Since end users will experience some disruption of service and possibly a performance degradation during failover, they might become concerned about the availability and stability of their applications. When the cluster is initially brought into a production environment, it might be beneficial to describe in a memorandum the effects a cluster will have on the end users' information systems needs.

When a failover or failback event occurs, end users will not be able to access their application and/or data. When users have been properly forewarned of the effects of operating in a clustered environment, they will more readily recognize when a failover or failback event is occurring or has occurred. Most will wait a few seconds to several minutes before trying to reconnect to their application and/or data.

Remotely Managing a Cluster

Both Compaq Insight Manager 7 and Microsoft Cluster Administrator can be run from network client machines. Each program allows you to monitor multiple clusters from a single remote client. Refer to the corresponding documentation for these products to determine how to set up and configure each of these programs to run remotely.

Cluster Events

Cluster events are logged in the Microsoft operating system Event Log. Use the Microsoft operating system Event Viewer to view the data in the Event Log.

Compaq Insight Manager 7

Compaq Insight Manager 7 is the strategic direction for Compaq to provide systems management and administration from the Web. Incorporated within Compaq Insight Manager 7 is Cluster Monitor, a real-time cluster monitoring system for ProLiant clusters using the Microsoft operating system with Microsoft Cluster software. The combination of Compaq Insight Manager 7 and Cluster Monitor provides complete systems and cluster monitoring and management.

Compaq Insight Manager 7 extends management beyond a single system to management of multiple systems using a standard Web browser as the user interface. By enabling browser access to both the managed device and the management application, Compaq Insight Manager 7 provides management of devices and groups of devices anywhere, anytime. Additionally, Compaq Insight Manager 7 provides real-time management access to Web-enhanced devices and proxy management of traditional SNMP and DMI V2 systems.

Compaq management agents provide health status to Compaq Insight Manager 7. The agents interpret data supplied by the device drivers into useful information that assists the user in correctly diagnosing the problem. Compaq Insight Manager 7 then provides extensive help information on how to address the problem, shortening the resolution time.

Compaq Insight Manager 7 offers a simple, industry-standard approach to management for all managed devices. This provides a common data repository, a consistent user interface, and the ability to correlate events providing intelligent information about the infrastructure at the department or domain level.

Compaq Insight Manager 7 provides device status summary of all managed devices. For Compaq systems this is a summary of the health status of all the subsystems. The Compaq management agents provide real-time access to the details found in the subsystem health status information.

Compaq Insight Manager 7 offers secure access to the management application and Web-enhanced management agents. This enhances the security provided by traditional SNMP agents. The management application and management agents authorize the user to look at information, change information, and execute operations.

Compaq Insight Manager 7 generates subsystem and component views through direct access to networked devices by means of a standard Web browser, offering customers a quick and convenient way to diagnose system status.

Compaq Insight Manager 7 and Cluster Monitor helps you focus on your computing environment from the perspective of Microsoft Cluster software clusters and their attributes.

Cluster Monitor

Cluster Monitor is a Web-based monitoring subsystem for Compaq Insight Manager 7. With Cluster Monitor, you can view all clusters from a single browser and configure monitor points and specific operational performance thresholds that will alert you when these thresholds have been met or exceeded on your application systems. Cluster Monitor relies heavily on the Compaq Insight Manager 7 Web-enabled agents for basic information about system health. It also has custom agents that are designed specifically for cluster health. Cluster Monitor provides access to the Compaq Insight Manager 7 alarm, device, and configuration information.

Cluster Monitor has three distinct informational areas that can be used to meet individual operational needs:

- A problem window with a prioritized cluster event list sorted by severity for the clusters that are under the administrator's control

- A topology tree displaying all clusters and their respective monitor points, including Microsoft Cluster software health, processor, bus, disk, or network usage and performance thresholds

- A detailed problem definition based on monitored conditions and a proposed resolution to the problem, if one can be determined, with drill-down ability to the specific device or system causing a negative or unstable cluster state and the ability to perform corrective actions

Cluster Monitor management information includes:

- Ownership

- Identified cluster resources

- System hardware and software configuration

- Emergency contacts in the event of cluster state changes

- Application software installed

Use Cluster Monitor to:

- Stay informed of changes in cluster status through cluster alert notifications.

- Monitor cluster status by viewing a list of cluster alerts.

- Investigate the sources of specific alerts.

- Browse cluster and component status in a tree hierarchy.

- Display utilization or status data for specific cluster attributes.

- Create Compaq Insight Manager 7 notifications to escalate or document cluster problems.

Cluster Monitor supports these attributes:

- Disk space

- CPU utilization

- MSCS cluster status

- Node Environment (Compaq Management Agent) status

Cluster Monitor uses pop-up notifications, alerts in the alert list, colored icons in a tree hierarchy, and attribute-specific data displays to reveal the state of clusters, their nodes, and attributes.

Refer to the Compaq Insight Manager 7-Cluster Monitor documentation for specific information about how to use Compaq Insight Manager 7 and Cluster Monitor.

Compaq StorageWorks Command Console

The Compaq StorageWorks Command Console (SWCC) is a graphical user interface (GUI) used to create and manage storage subsystems. While the CLI window provides very detailed control over the storage subsystem, the SWCC replicates most of the CLI window functions in a graphic form and provides a user-friendly method of executing CLI commands. Refer to the SWCC documentation for more information.

Secure Path Manager

Secure Path Manager is a component of the Secure Path software included in the ProLiant Cluster HA/F500 Enhanced Cluster Kit. Secure Path Manager is the client application used to manage multipath storage subsystem configurations. It displays a graphical representation of the current multipath environment and indicates the location and state of all configured storagesets on each of the paths. To facilitate static load balancing, Secure Path Manager provides the capability to move storagesets between paths.

Microsoft Cluster Administrator

Microsoft Cluster Administrator allows you to manage the groups, resources, and operating state of the cluster. It gives you the ability to:

- View the current status of cluster groups and resources.

- Bring groups and resources online and offline.

- Manually move groups and resources to another node.

- Manually move groups and resources to their preferred server.

- Pause groups and resources.

- Restructure a group resource dependency tree.

Microsoft Cluster Administrator can be run remotely or on a cluster node. If Microsoft Cluster Administrator is installed remotely, the remote node must be in the same domain as the cluster node. When running Microsoft Cluster Administrator remotely, you should connect to the cluster by using the cluster IP address or cluster name.

Refer to the *Microsoft Cluster Server Administrator Guide* for a thorough description of Microsoft Cluster Administrator features and specifics on how to use the utility.

Modifying Physical Cluster Resources

The following sections provide information on modifying the physical cluster resources.

Removing Shared Storage

Each cluster node, and the cluster as a whole, depends on the shared storage for data, log, and possibly application program files. If the shared storage is removed, all clustered applications that are dependent on that shared storage system will be offline. Network clients will not have access to the clustered applications.

Before removing a shared storage system, use Microsoft Cluster Administrator to bring offline all cluster groups and resources on both cluster nodes. This action ensures that the groups and resources are gracefully brought offline, rather than causing an abrupt interruption of service by powering off the storage subsystem.

Adding and Configuring Physical Hard Drives

The following sections provide step-by-step instructions for adding a single physical disk to an HA/F500 cluster while keeping the cluster online at all times.

If you need more information on Microsoft Cluster software, refer to the documentation included on the Windows NT Server 4.0, Enterprise Edition CD, or the Windows 2000 CD.

1. Shut down and power off one of the cluster servers (for example: Node B).

2. If you are installing new hard drives, insert them into the external storage unit at this time. Allow the new hard drives to spin up.

3. Open the **HSG80 Storage** window. If you have installed a new hard drive, then get the storage system to recognize that new drive.

4. Create the new logical unit. After creating the logical unit, exit the utility.

5. Power up Node B. Log on to an account with administrator rights.

6. Wait for Node B to join the cluster.

7. Shut down and power off Node A. Wait for the cluster resources to failover to Node B.

Alternate Configuration of Hard Disks

1. Open **Hyperterminal.**

2. Connect to the storage unit.

3. Enable access for the new disks on their given port.

Configuring a New NTFS Volume

1. Open **Disk Administrator** or **Disk Management** for Windows 2000. A pop-up window displays informing you that a new disk has been added. Click **OK** on the pop-up window to update the system configuration.

2. A pop-up window displays asking if you want to assign a disk signature to the new disk. Click **Yes** for the newly created logical volume.

   NOTE: For Windows 2000 do not upgrade the volume to a dynamic disk. Leave the volume as a basic disk. Dynamic disks are not supported for cluster shared storage disks.

3. The newly added disk displays as free space in the **Disk Administrator** or **Disk Management** window. Create one partition on the entire disk volume.

   NOTE: Do not sub-partition the logical volume using Windows Disk Administrator or Disk Management for Windows 2000.

4. Format the new partition using NTFS.

5. Assign a drive letter for the new partition.

6. Close **Disk Administrator** or **Disk Management.**

Adding a Physical Disk Resource to the Cluster

To add a disk resource to the cluster:

1. Open **Cluster Administrator** and connect to the cluster.

2. If you want to add the disk resource to an existing group, then skip to step 5. Otherwise, select **File, New, Group** from the menu bar.

3. Enter a name (description is optional) for the new group. Click **Next.**

4. Select the preferred nodes and preference order for this group. Click **Finish.** Click **OK** after creating the new group.

5. Select **File, New, Resource.**

6. Enter a name (description is optional) for the new resource.

7. For **Resource Type,** select **Physical Disk.**

8. Select the group to which you want to add the disk resource. Click **Next.**

9. Accept the default nodes for which the resource may be brought online. Click **Next.**

10. Specify which resources you want the new disk resource to be dependent on. Click **Next.**

11. Select the disk volume you want to assign to the disk resource. Click **Finish.** Click **OK** after creating the new resource.

12. Bring the new physical disk resource online by right-clicking its icon and selecting **Bring Online.**

13. Power up Node A and log on to the Microsoft operating system.

Physically Replacing a Cluster Node

At some time you may need to permanently replace one of the nodes in the cluster. For example, over time you may slowly increase the number of applications running on the cluster, which may require you to replace an existing node with a new, more powerful and expandable cluster node.

To simplify integration of the new cluster node, Compaq recommends that you retain the Microsoft operating system boot drives from the node being replaced for use in the replacement node. If this is not possible, you need to perform the steps in the "Installing a New Microsoft Operating System Boot Drive" section later in this chapter.

Replacing a Cluster Node

To replace an existing cluster node:

1. Fail over to the remaining node all cluster groups that are running on the node being replaced, if not already accomplished. During replacement of the node, all cluster groups must be running on the remaining cluster node.

   NOTE: Certain cluster-aware applications do not allow replacements of the node on which the application was first installed. Remove these applications from the cluster before replacing the first node the application was installed on.

2. Evict the node to be replaced from Microsoft Cluster Administrator.

3. Shut down and power off the node that is being replaced.

4. Remove the interconnect, LAN, Fibre Channel, and power cables. Remove any other cables that are attached to the node. If the node is in a rack, remove it from the rack.

5. For all hardware devices that will be used in the new node, remove them from the node being replaced and place them in the new node. Install all other hardware devices into the new node.

6. If the new node is part of a rack system, place the server in the rack. Attach the interconnect, LAN, Fibre Channel, and power cables.

If you are using the Microsoft operating system boot drives from the replaced node in the new node, power on the new node and follow the steps described in the "Verifying Integration of the Replacement Node" section later in this chapter.

Installing a New Microsoft Operating System Boot Drive

A new Microsoft operating system boot drive requires installation of the Microsoft operating system, configuration of the networking components of the new node, and installation of Microsoft cluster software.

1. Boot the disk from the SmartStart CD.

2. Select **System Erase Utility.**

3. When the System Erase Utility finishes, select **Assisted Integration.**

4. Install Microsoft cluster software. Be sure you join an existing cluster.

    When Microsoft cluster software is installed, registry replication from the original node to the replacement server occurs automatically. All cluster resources, groups, and dependency trees will be available on the replacement server.

5. Configure the network devices for the new node according to Microsoft cluster software recommendations. The new node must be in the same domain as the existing node. If a dedicated interconnect is used, then the interconnect controller for the new node must be on the IP network as the corresponding controller in the other cluster node.

6. Shut down the replacement server, then reboot it and allow it to load the Microsoft operating system and Microsoft cluster software.

Verifying Integration of the Replacement Node

After you have installed the replacement node and the new Microsoft operating system boot drive (if necessary), verify integration of the replacement node into the cluster.

Start Microsoft Cluster Administrator and verify that all groups, resources, and nodes are visible. When applicable, fail back applications and resources from the original cluster node to the newly integrated node.

At this point, the cluster should be running in a normal state.

If you encounter problems, the first troubleshooting activity should be to bring all applications offline on each node and shut down the servers. Restart the servers and allow the Microsoft operating system and Microsoft cluster software to start. Use Microsoft Cluster Administrator to verify the cluster status. If you still encounter problems, refer to Chapter 6.

Cluster Backup

Cluster backup is essential for both business-critical and nonbusiness-critical data. As data becomes an increasingly valued company asset, it needs to be backed up on a regular basis. The process of backing up data ensures that company assets are secure and available when a disaster strikes. The cluster itself provides a high degree of application availability but does not prevent a user from deleting or corrupting a file or set of files. Backing up cluster data increases the overall level of data availability.

Several methods for backing up clustered systems exist:

- Attaching both cluster nodes to a single tape backup device, such as the Fibre Channel-based Compaq StorageWorks Enterprise Backup Solution

- Attaching two-tape backup subsystems, where one tape backup subsystem is attached to each cluster node

- Establishing a separate backup server that uses the standard public client LAN as the backup path

Tape backup software varies in its level of cluster-aware integration and operation. Contact your tape backup software vendor for more details regarding their specific cluster implementation and support.

Upgrading the ACS Version

To upgrade the ACS version on a ProLiant Cluster with a Windows 2000 operating system, perform a shutdown upgrade. Do **not** perform a rolling upgrade; all servers will need to be shut down during the upgrade process.

Each Windows 2000 cluster node must be rebooted independently from each other so that the proper drive discovery can occur after the upgrade.

Follow the procedures included with the ACS platform kit.

6

Troubleshooting the ProLiant Cluster HA/F500

This chapter addresses potential problems you may encounter as you install, configure, test, and operate the ProLiant Cluster HA/F500.

Verifying Cluster Configuration

The following sections provide information on verifying the cluster configuration.

Verifying Device Driver Initialization

A problem with the Windows NT Hardware Abstraction Layer (HAL) may prevent the host bus adapter device driver (LP6NDS35.SYS) from initializing during system boot. The consequences of this problem are that none of the devices connected to the host bus adapter will be available and an event log entry will be entered in the system event log.

The problem is caused by the Windows NT HAL reconfiguring the PCI adapter with conflicting resources. The following workaround causes the HAL to use the BIOS-assigned defaults and not reassign PCI resources.

1. Use Notepad or Edit to edit the *BOOT.INI* file as follows:

   a. Remove the read-only file attribute.
      Example: `attrib –r –h -s c:\boot.ini`

   b. Add the /PCILOCK option to the system boot entry.

   ```
   (boot loader)
   timeout=30
   default=multi(0)disk(0)rdisk(0)partition(2)\WINNT
   (operating systems)
   multi(0)disk(0)rdisk(0)partition(2)\WINNT="Windows NT
   Server,
   Enterprise Edition Version 4.00"/pcilock
   ```

2. Reboot the server.

Verifying the Windows NT Registry Device Driver Parameters

For proper operation of the storage subsystem configured in a Microsoft operating system cluster, the driver parameters in the operating system registry must be set correctly.

To verify the operating system registry settings:

1. Run the Fibre Channel Software Setup Utility.

2. Select the correct Fibre Channel system (multibus or transparent, hub or switch) and select **Custom Install.**

3. Install **Extended Configuration** if more than two clusters access the same HSG80 storage system. This action automatically corrects the registry settings configured in the system.

4. Reboot the system after completing the Fibre Channel Software Setup.

Verifying Connectivity to the Fibre Channel Arbitrated Loop

Verify that the storage subsystem and the host servers are properly connected to the Fibre Channel Arbitrated Loop.

Storage Subsystem Connectivity to the Loop

Enter the CLI command `show this_controller` to ensure that the storage controller is connected to the loop. If the port 1 topology is in an offline state, refer to the HSG80 documentation for more information.

The port 1 topology should be LOOP_HARD and indicate that the port state is Loop Up or Standby. If one storage controller indicates a Loop Up state, the other storage controller will be in a Standby mode.

Host Server Connectivity to the Loop

To verify that the Microsoft operating system servers are connected to the loop, enter the CLI command `show connection`. The output of this command shows status information about present and previous connections to the loop. The **Unit Address** column refers to the Arbitrated Loop Physical Address. The **Status** column indicates the port is online to the other storage controller.

It is recommended that you rename the connection-name field to the computer name of the server. This action facilitates the ability to discern loop connections.

Troubleshooting Installation Problems

The following sections provide information on troubleshooting the ProLiant Cluster HA/F500. For troubleshooting installation problems with the Microsoft operating system or Microsoft cluster software, refer to the Microsoft website at http://support.microsoft.com.

Also refer to the Microsoft white paper *Microsoft Cluster Server Troubleshooting and Maintenance.*

Configuring Multiple Storagesets Error

When configuring multiple storagesets, the following error may occur:

```
Error 8020: No metadata found on container. An INITIALIZE
DISKxxxxx must be issued before this container is used
```

To correct this error:

1. Allow enough time for the previous storagesets to initialize before continuing.

2. Recreate the storageset that erred and then continue configuring the remaining storagesets.

Discovery of Units Error

After configuring the storagesets for the first time, Secure Path may encounter an error during the discovery of the units when the server is restarted. The second server is powered off at this time. An Administrator Alert, "Unable to Configure External Storage," is received.

To correct the error, reboot the server as many times as needed to discover all the units.

Physical Disk Resources Cannot Be Brought Online

- For Windows NT, both cluster nodes should be rebooted after installing Microsoft Cluster software. If you have not done so, reboot the two nodes.

- Verify that there are no hardware errors or transport problems by using Event Viewer. Look in the Event Log for disk I/O error messages or indications of problems with the communications transport.

- Microsoft Cluster Administrator takes a snapshot of the registry when it starts up. However, it may take up to a minute after the second node is rebooted for the disk signatures to be written to both registries. You may not have waited long enough to view the most up-to-date data. Wait a minute, then click **Refresh** (or press the **F5** key).

Amber Light Flashes on Switch Port

If an amber light flashes while verifying access during switch installation, restart the switch and begin again.

Troubleshooting Node-to-Node Connectivity Problems

The following sections provide node-to-node connectivity troubleshooting.

Second Node Cannot Join the Cluster

- Check the cluster IP address resource properties to ensure that the cluster has a valid IP address and subnet mask, and that the IP address does not conflict with an IP address that is already in use on the network.

- If the cluster nodes use DHCP to obtain noncluster IP addresses, use ipconfig.exe from the command prompt to verify that you have a valid primary IP address for the adapters in question. If the second IP address listed is 0.0.0.0, then the primary address is invalid.

Troubleshooting Shared Storage Problems

This section addresses potential problems arising from the use of the StorageWorks storage subsystem as a shared storage device in a cluster. This section does not address problems that are specific to the storage subsystem itself or to the use of the storage system in a stand-alone server configuration. For those issues, refer to the documentation provided with the storage subsystem.

Nodes Cannot Connect to the Shared Drives

- Verify that a physical connection exists from the second node to the Fibre Channel storage hub or switch:

  — Verify that the shortwave Gigabit Interface Converters (GBICs) are properly seated.

  — Verify that all Fibre Channel cables are properly connected to their GBICs.

- Verify that the shared drives are assigned the same drive letters on both nodes. To do so, run Windows Disk Administrator on each node and verify that all shared drives have been assigned identical drive letters. The drive letter assignments must be permanent.

- If the second node was powered on before either the Fibre Channel storage hub or switch or the storage subsystem; shut down the second node, turn it off, then turn it on. The storage subsystem components must be powered on before the servers.

Troubleshooting Cluster Group and Cluster Resource Problems

Cluster group and cluster resource problems are related to Microsoft Cluster Administrator. For troubleshooting tips on this topic, see the *Microsoft Cluster Server Administrator Guide* and Microsoft Cluster Administrator online documentation.

A

Versions of Referenced Software

The following table is a list of all software and firmware versions mentioned in this guide.

IMPORTANT: Refer to the HA/F500 order and configuration guide on the Compaq High Availability website (www.compaq.com/highavailability) for information about software and firmware updates recommended or required for the ProLiant cluster.

Table A-1: Referenced Software

Software Package	Version Number
SmartStart CD	5.40 or later
Compaq Support Paq for Windows	5.40A or later
SANworks Secure Path Version 4.0 for Windows	4.0
Compaq Insight Manager 7	3.0 or later
ACS Controller Software	8.6F or 8.7F

continued

Table A-1: Referenced Software *continued*

Software Package	Version Number
KGPSA-BC Host Bus Adapter Firmware	3.20x7
KGPSA-CB Host Bus Adapter Firmware	3.82a1/1.60a4
FCA 2101 Host Bus Adapter Firmware	3.82a1/1.60a5
Platform kits	Compaq Array Controller Software Solution v8.6-4 Solution kit or
	Compaq Array Controller Software v8.6+ Solution kit or
	Compaq Array Controller Software v8.7 Solution kit
Microsoft Windows NT Server, Enterprise Edition	4.0
Service Pack	6a
Microsoft Windows 2000 Advanced Server	5.0
Service Pack	2 or later
Microsoft Windows 2000 Datacenter Server	
Service Pack	2 or later

Glossary

ACS
See array controller software.

active/active
A dual controller, dual adapter, storage subsystem configuration in which both controller-adapter I/O paths have access to separate logical units (LUNs). *Also known as* multibus.

active/standby
A dual controller, single adapter storage subsystem configuration in which one controller is in an online state and has control of the logical storage units. The other controller is in a standby state. *Also known as* transparent.

adapter
A device that converts the protocol and hardware interface of one bus type into another without changing the function of the bus.

ALPA
See arbitrated loop physical address.

arbitrated loop
A loop type of topology where two or more ports can be interconnected, but only two ports at a time can communicate.

arbitrated loop physical address (ALPA)
A one-byte value used to identify a port in an arbitrated loop topology.

array controller software (ACS)
Software contained on a removable ROM program card that provides the operating system for the array controller.

availability
A measure of how well a computer system or cluster can continuously deliver services to its clients. Availability is typically expressed as a percentage, with 100 percent being the best possible rating.

CLI
See command line interpreter.

cluster
A group of systems that work collectively as a single system to provide fast, uninterrupted computing service. Clustering is a way to increase availability, processing capacity, and I/O bandwidth.

cluster group
A collection of interdependent resources that logically represents a clustered client/server function. This is a user-definable entity used by Microsoft Cluster Server software.

command line interpreter (CLI)
The configuration interface used to operate the controller software.

container
1. Any entity that is capable of storing data, whether it is a physical device or a group of physical devices.

2. A virtual, internal controller structure representing either a single disk or a group of disk drives linked as a storageset.

controller
A hardware device that, with proprietary software, facilitates communications between a host and one or more devices organized in an array.

driver
A hardware device or a program that controls or regulates another device. For example, a device driver is a driver developed for a specific device that allows a computer to operate with the device, such as a printer or a disk drive.

dual-redundant configuration
A controller configuration consisting of two active controllers operating as a single controller. If one controller fails, the other controller assumes control of the failing controller devices.

Ethernet
A standard network protocol that operates mostly on a physical level, using network interface cards and cabling to transmit data between computers. Transfer rates are normally 10 or 100 megabits per second.

fabric
Multiple Fibre Channel switches interconnected and using Fibre Channel methodology for linking nodes and routing frames in a Fibre Channel network.

failback
1. The process that takes place when a previously failed controller is repaired or replaced and reassumes the workload from a companion controller.

2. The process that takes place when the operation of a previously failed cluster group moves from one cluster node back to its primary node.

failover
1. The process that takes place when one controller in a dual-redundant configuration assumes the workload of a failed companion controller. Failover continues until the failed controller is repaired or replaced.

2. The process that takes place when the operation of a cluster group moves from one cluster node to another node in the same cluster.

FC-AL
The Fibre Channel Arbitrated Loop standard. *See* arbitrated loop. *See* Fibre Channel.

FC-SW
Fibre Channel Switched Topology and Switch Controls.

Fibre Channel
An IEEE standard for providing high-speed data transfer among workstations, mainframes, supercomputers, desktop computers, storage devices, and display devices.

heartbeat
A signal transmitted between cluster nodes to indicate whether the nodes are operating.

high availability
A term used to identify a computer system that can continuously deliver services to its clients.

host
The primary or controlling computer to which a storage subsystem is attached.

host bus adapter
A device that connects a host system to a SCSI bus. The host bus adapter usually performs the lowest layers of the SCSI protocol. This function may be logically and physically integrated into the host system.

interconnect
A physical connection between cluster nodes that transmits intracluster communication.

I/O
A term that refers to input and output functions.

IP address
Internet Protocol address. An address assigned to a network interface card, which computer entities use to locate and communicate with each other. IP addresses can be statically or dynamically assigned.

load balancing
See static load balancing.

local terminal
A terminal plugged into the EIA-423 maintenance port located on the front bezel of the controller. *See also* maintenance terminal.

logical unit (LUN)
A physical or virtual device addressable through a target ID number. LUNs use the target bus connection to communicate on the SCSI bus.

logical unit number

1. A value that identifies a specific logical unit belonging to a SCSI target ID number.

2. A number associated with a physical device unit during the I/O operations of a task. Each task in the system must establish its own correspondence between logical unit numbers and physical devices.

loop

See arbitrated loop.

LUN

See logical unit (LUN).

maintenance terminal

An EIA-423-compatible terminal used with the controller. This terminal is used to identify the controller, enable host paths, enter configuration information, and check the controller status. *See also* local terminal.

multibus

A dual controller, dual adapter, storage subsystem configuration in which both controller-adapter I/O paths have access to separate logical units. *Also known as* active/active.

NIC

Network Interface Controller. A board that allows a computer to be connected with a network and that works with the network operating system to control the flow of information over the network.

node

An individual server in a cluster.

NTFS

NT File System. A file organization system by which data is stored and accessed in a Windows NT operating system.

other controller

In a dual-controller configuration, the storage controller that is not connected to the maintenance terminal cable. *See also* this controller.

partition
A logical division of a container, represented to the host as a logical unit.

PCMCIA
Personal Computer Memory Card Industry Association. An international association formed to promote a common standard for PC card-based peripherals to be plugged into notebook computers. The card, commonly known as a PCMCIA card, is about the size of a credit card.

port
1. In general terms, a logical channel in a communication system.

2. The hardware and software used to connect a host controller to a communications bus, such as a SCSI bus or serial bus.

program card
The PCMCIA card containing operating software for the storage controller.

RAID
Redundant Array of Inexpensive Disks. A method of using hard disk drives in an array to provide data redundancy to increase system reliability and performance.

redundancy
The provision of multiple interchangeable components to perform a single function in order to cope with failures and errors.

reliability
The continuous integrity of a system (server, storage, network, or cluster).

resource
A software or hardware entity upon which a client/server application or service is dependent. As it pertains to Microsoft Cluster Server, a cluster resource must have the ability to be managed by the cluster and must reside on one of the cluster nodes. A resource can be a member of only one group.

SCSI
Small Computer Systems Interface. A standard parallel interface for rapid data transmission.

ServerNet
A bidirectional, high-bandwidth, low-latency, redundant path network interconnect.

shared resource

A type of cluster organization in which some resources are accessible to all systems in the cluster.

static load balancing

A method for manually balancing the amount of work undertaken by each storage controller in a single storage subsystem.

storageset

A group of devices configured with RAID techniques to operate as a single container.

storage subsystem

The controllers, storage devices, shelves, cables, and power supplies used to form a mass storage subsystem.

system

A complete computer system capable of operating independently.

this controller

The storage controller to which the maintenance terminal cable is physically connected. *See also* other controller.

transparent

A dual-controller, single-adapter storage subsystem configuration in which one controller is in an online state and has control of the logical storage units. The other controller is in a standby state. *Also known as* active/standby.

virtual disk

See logical unit.

Index